arXiv:1703.05038v1 [math.NA] 15 Mar 2017

Harmonic Mean Iteratively Reweighted Least Squares for Low-Rank Matrix Recovery

Christian Kümmerle*, Juliane Sigl†

March 16, 2017

Abstract

We propose a new iteratively reweighted least squares (IRLS) algorithm for the recovery of a matrix $X \in \mathbb{C}^{d_1 \times d_2}$ of rank $r \ll \min(d_1, d_2)$ from incomplete linear observations, solving a sequence of low complexity linear problems. The easily implementable algorithm, which we call harmonic mean iteratively reweighted least squares (`HM-IRLS`), optimizes a non-convex Schatten-p quasi-norm penalization to promote low-rankness and carries three major strengths, in particular for the matrix completion setting. First, the algorithm *converges globally* to the low-rank matrix for relevant, interesting cases, for which any other (non-)convex state-of-the-art optimization approach fails the recovery. Secondly, `HM-IRLS` exhibits an empirical recovery probability close to 100% even for a number of measurements very close to the theoretical lower bound $r(d_1 + d_2 - r)$, i.e., already for *significantly fewer linear observations* than any other tractable approach in the literature. Thirdly, `HM-IRLS` exhibits a *locally superlinear rate of convergence (of order* $2 - p$*)* if the linear observations fulfill a suitable null space property. While for the first two properties we have so far only strong empirical evidence, we prove the third property as our main theoretical result.

1 Introduction

The problem of recovering a low-rank matrix from incomplete linear measurements or observations has gained considerable attention in the last few years due to the omnipresence of low-rank models in different areas of science and applied mathematics. Low-rank models arise in a variety of areas such as system identification [35, 36], signal processing [1], quantum tomography [27, 25] and phase retrieval [9, 5, 26]. An instance of this problem of particular importance, e.g., in recommender systems [46, 24, 8], is the *matrix completion* problem, where the measurements correspond to entries of the matrix to be recovered.

Although the low-rank matrix recovery problem is NP-hard in general, several tractable algorithms have been proposed that allow for provable recovery in many important cases. The *nuclear norm minimization* (`NNM`) approach [16, 8], which solves a surrogate semidefinite program, is particularly well-understood. For `NNM`, recovery guarantees have been shown for a number of measurements on the order of the information theoretical lower bound $r(d_1 + d_2 - r)$, if r denotes the rank of a $d_1 \times d_2$-matrix [43, 8]; i.e., for a number of measurements $m \geqslant \rho r(d_1 + d_2 - r)$ with some oversampling constant $\rho \geqslant 1$. Even though `NNM` is solvable in polynomial time, it can be computationally very demanding if the problem dimensions are large, which is the case in many potential applications. Another issue

*Department of Mathematics, Technische Universität München, Boltzmannstr. 3, 85748 Garching/Munich, Germany
E-mail: `christian.kuemmerle@ma.tum.de`

†Department of Mathematics, Technische Universität München, Boltzmannstr. 3, 85748 Garching/Munich, Germany
E-mail: `juliane.sigl@ma.tum.de`

is that although the number of measurements necessary for successful recovery by nuclear norm minimization is of *optimal order*, it is not *optimal*. More precisely, it turns out that the oversampling factor ρ of nuclear norm minimization *has to be much larger than the oversampling factor of some other, non-convex algorithmic approaches* [59, 50].

These limitations of convex relaxation approaches have led to a rapidly growing line of research discussing the advantages of non-convex optimization for the low-rank matrix recovery problem [30, 50, 28, 31, 58, 51, 54, 57, 56, 52, 60]. For several of these non-convex algorithmic approaches, recovery guarantees comparable to those of `NNM` have been derived [6, 52, 48, 60, 49]. Their advantage is a higher empirical recovery rate and an often more efficient implementation, but their success usually relies on a distinct, possibly expensive initialization step.

1.1 Contribution of this paper

In this spirit, we propose a new iteratively reweighted least squares (IRLS) algorithm for the low-rank matrix recovery problem that strives to minimize a non-convex objective function based on the Schatten-p quasi-norm

$$\min \|X\|_{S_p}^p \text{ subject to } \Phi(X) = Y,$$

for $0 < p < 1$, where $\Phi : \mathbb{R}^{d_1 \times d_2} \to \mathbb{R}^m$ is the linear measurement operator and $Y \in \mathbb{R}^m$ is the data vector defining the problem. The overall strategy of the proposed IRLS algorithm is to mimic this minimization by a sequence of weighted least squares problems. This strategy is shared by the related previous algorithms [19, 39] proposed already some time ago. However, as an important innovation, we propose the use of a different type of weight matrices, so-called *harmonic mean weight matrices*, which can be interpreted as the harmonic mean of left- and right-sided weight matrices introduced in [19, 39]. This motivates the name harmonic mean iteratively reweighted least squares (`HM-IRLS`) for the algorithm.

This new choice of weight matrices is *more symmetrical* than the weight matrices proposed in [19, 39], and enables `HM-IRLS` to use the information in both the column and the row space of the iterates, unlike in the mentioned, previous IRLS algorithms. More specifically, the structure of harmonic mean weight matrices is able to achieve a better alignment of the left-singular and right-singular vectors of the iterates with the ones of the low-rank matrix to be recovered in the non-convex case, i.e., if $p < 1$.

To put our work into the proper context, we note that the authors of [10, 13] proposed IRLS algorithms for the related sparse vector recovery problem [21, 17], which strive to minimize the ℓ_p quasi-norm under a linear constaint for $0 < p \leqslant 1$. A thorough analysis for these IRLS algorithm has been provided in [13] and in particular, for $p < 1$, a locally superlinear rate of convergence of order $2 - p$ was proven and also verified very accurately in numerical experiments (see Figure 4.8 in [13]). This work can also be seen as the inspiration of the papers [19, 39] about IRLS algorithms for the low-rank matrix recovery problem. Unfortunately, the algorithms considered in these papers are not able to generalize all the properties of [13]. Importantly, the theory of [19, 39] does not predict a superlinear convergence rate for non-convex penalizations corresponding to $p < 1$ and neither was it observed numerically.

In this paper, we show that a choice of harmonic mean weight matrices as in `HM-IRLS` is able to overcome the shortcomings of the weight matrices of [19, 39] for non-convex penalizations corresponding to $p < 1$. This gives rise to an efficient algorithm for the low-rank matrix recovery problem with fast convergence and superior performance in terms of sample complexity, also compared to algorithms based on strategies different from IRLS.

In particular, we extend the theoretical guarantees of previous low-rank IRLS algorithms, both general ones and those based on a Schatten-p *null space property* (NSP) of the measurement operator [40, 21], to `HM-IRLS`. More precisely, we prove convergence of accumulations points of the sequence of iterates to stationary points of a smoothed Schatten-p functional, under the linear constraint.

Our main theoretical result is that HM-IRLS exhibits a locally superlinear convergence rate of order $2-p$ in the neighborhood of a low-rank matrix for the non-convexity parameter $0 < p < 1$ connected to the Schatten-p quasinorm, if the measurement operator fulfills the mentioned NSP of sufficient order. For $p \ll 1$, this means that the convergence rate is *almost quadratic*.

Measurement models that fulfill the Schatten-p NSP of the required order with high probability for optimal sample complexities include random operators, e.g., with i.i.d. Gaussian entries [7]. It is important to mention that the measurement operator in the matrix completion setting, i.e., the choice of the measurement operator $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$, where we are given m sample entries $\Phi(X)_\ell = X_{i_\ell, j_\ell}$, for $\ell = 1, \dots, m$ and some $i_\ell \in [d_1]$ and $j_\ell \in [d_2]$ depending on ℓ, does not fulfill a null space property and hence these guarantees do not apply directly.

Nevertheless, due to the popularity of the problem and for reasons of comparability with other algorithms, we conduct numerical experiments to explore the empirical performance of HM-IRLS mainly for the matrix completion setting. Surprisingly enough we observe that the theoretical results comply with our numerical experiments also for matrix completion. In particular, the theoretically predicted local convergence rate of order $2 - p$ can be observed very precisely in this important measurement model as well (see Figures 1-3).

As already mentioned, this local superlinear convergence rate *is unprecedented* in the IRLS variants of [19, 39]. Moreover, other ways of defining a more symmetrical reweighting, such as using the arithmetic mean of left and right weight matrices, do not lead to superlinear convergence rates: neither can a superlinear rate be verified numerically, nor is it possible to show such a rate by our proof techniques. *Hence, we consider the HM-IRLS as the presently best extension of [13] to the low-rank matrix recovery problem.* To the best of our knowledge, HM-IRLS is the first algorithm for low-rank matrix recovery which achieves superlinear rate of convergence for low complexity measurements as well as for larger problems. In particular, although the algorithms appearing in [38] and NewtonSLRA in [45] do show superlinear convergence rates, the first is not competitive to HM-IRLS in terms of sample complexity and the second has neither applicable theoretical guarantees for most of the interesting probems nor the ability of solving medium size problems.

Additionally, we conduct extensive numerical experiments comparing the efficiency of HM-IRLS with previous IRLS algorithms [19, 39], Riemannian optimization techniques [54], alternating minimization approaches [28, 51], algorithms based on iterative hard thresholding [34, 4], and others [41], in terms of sample complexity, again for the important case of *matrix completion*. More precisely, the experiments are executed to compare the empirical recovery probabilities of the algorithms for number of measurements $m := \rho r(d_1 + d_2 - r)$, given problem parameters r, d_1, d_2 characterizing the randomly sampled ground truth matrices $X_0 \in \mathbb{R}^{d_1 \times d_2}$ of rank r and varying oversampling factors $\rho \geqslant 1$.

The experiments lead to the following observation: HM-IRLS recovers low-rank matrices systematically with an optimal number of measurements, in the sense that an oversampling factor $\rho \approx 1$ that is very close to 1 suffices for recovery with high empirical probability, whereas all the other algorithms described above necessitate a *significantly larger* factor ρ to be able to recover X_0. We consider this result to be remarkable, as it means that for problems of moderate dimensionality (matrices of $\approx 10^7$ variables, e.g. $(d_1 \times d_2)$-matrices with $d_1 \approx d_2 \approx 5 \cdot 10^3$) *the proposed algorithm needs fewer measurements for the recovery of a low rank matrix than all the state-of-the-art algorithms we included in our experiments.*

An important practical advantage of HM-IRLS is that it does not depend, unlike many non-convex methods different from IRLS such as alternating minimization [28, 31] and related methods [51, 60], on a distinct, well-chosen initialization. Instead, HM-IRLS can be used as a stand-alone algorithm to recover low-rank matrices starting from a trivial initialization. Interestingly, our numerical experiments suggest that HM-IRLS is able to *converge globally* to the low-rank matrix, even though it is based on an objective function which is highly non-convex. The global convergence is confirmed by the remarkbly high empirical recovery probability of HM-IRLS, which is close to 100% as soon as the oversampling factor ρ is raised sufficiently above 1. In contrast, other non-convex optimization

approaches fail as a consequence from a lack of either convergence or NSP stability guarantees. While a complete theoretical understanding of this behavior is not yet achieved, we regard the empirical evidence for the global convergence of `HM-IRLS` in a variety of interesting cases as overwhelming. In this context, we consider a proof of the global convergence of `HM-IRLS` for non-convex penalizations under appropriate assumptions as an interesting and very hard open problem.

1.2 Outline

We proceed in the paper as follows. In the next section, we provide some background on Kronecker and Hadamard products of matrices as these concepts are used in the analysis of the algorithm to be discussed. Moreover, we explain different reformulations of the Schatten-p quasi-norm in terms of weighted ℓ_2-norms, which lead to the derivation of the harmonic mean iteratively reweighted least squares (`HM-IRLS`) algorithm in section 3. We present our main theoretical results, the convergence guarantees and the locally superlinear convergence rate for the algorithm in section 4. Numerical experiments and comparisons to state-of-the-art methods for low-rank matrix recovery are carried out in section 5. In section 6, we interpret the algorithm's different steps as minimizations of an auxililary functional with respect to its arguments and show theoretical guarantees for `HM-IRLS` extending similar guarantees for related methods [19, 39]. After this, we detail the proof of the locally superlinear convergence rate under appropriate assumptions on the null space of the measurement operator.

2 Notation and background

2.1 General notation, Schatten-p and weighted norms

In this section, we explain some of the notation we use in the course of this paper.

Set set of matrices $X \in \mathbb{C}^{d_1 \times d_2}$ is denoted by $M_{d_1 \times d_2}$. Unless stated otherwise, vectors $x \in \mathbb{C}^d$ are considered as column vectors. We also use the vectorized form $X_{\mathrm{vec}} = \left[X_1^T, \dots, X_j^T, \dots, X_{d_2}^T\right]^T \in \mathbb{C}^{d_1 d_2}$ of a matrix $X \in M_{d_1 \times d_2}$ with columns X_j, $j \in \{1, \dots, d_2\}$. The reverse recast of a vector $x \in \mathbb{C}^{d_1 d_2}$ into a matrix of dimension $d_1 \times d_2$ is denoted by $x_{\mathrm{mat}(d_1,d_2)} = [X_1, \dots, X_j, \dots, X_{d_2}]$, where $X_j = [x_{(d_1-1)\cdot j+1}, \dots, x_{(d_1-1)\cdot j+d_1}]^T$, $j = 1, \dots, d_2$ are column vectors, or X_{mat} if the dimensions are clear from the context. Obviously, it holds that $X = (X_{\mathrm{vec}})_{\mathrm{mat}}$.

The identity matrix in dimension $d \times d$ is denoted by $\mathbf{I}_d$. With $\mathbf{0}_{d_1 \times d_2} \in M_{d_1 \times d_2}$ and $\mathbf{1}_{d_1 \times d_2} \in M_{d_1 \times d_2}$ we denote the matrices with only 0- or 1-entries respectively. The set of Hermitian matrices is denoted by $H_{d \times d} := \{X \in M_{d \times d} \mid X = X^*\}$. We write $X^+ \in M_{d_1 \times d_2}$ for the Moore-Penrose inverse of the matrix $X \in M_{d_1 \times d_2}$.

Let $U_d = \{U \in \mathbb{C}^{d \times d}; UU^* = \mathbf{I}_d\}$ denote the set of unitary matrices. We write the singular value decomposition of a matrix $X \in M_{d_1 \times d_2}$ as $X = U\Sigma V^*$ with $U \in U_{d_1}$ and $V \in U_{d_2}$ and $\Sigma \in M_{d_1 \times d_2}$, containing the singular values of X on the main diagonal. Let the entries of this diagonal of Σ be called $\sigma_i(X) = \Sigma_{ii} \geqslant 0$ for $i \in \{1, \dots, \min(d_1, d_2)\}$ and define the *Schatten-p (quasi-)norm* of $X \in M_{d_1 \times d_2}$ as

$$\|X\|_{S_p} := \begin{cases} \operatorname{rank}(X), & \text{for } p = 0, \\ \left[\sum_{j=1}^{\min(d_1,d_2)} \sigma_j^p(X)\right]^{1/p}, & \text{for } 0 < p < \infty, \\ \sigma_{\max}(X), & \text{for } p = \infty. \end{cases} \tag{1}$$

Note that for $p = 1$, the Schatten-p norm is also called *nuclear norm* and written as $\|X\|_* = \|X\|_{S_1}$. The *trace* $\operatorname{tr}[X]$ of a matrix $X \in M_{d_1 \times d_2}$ is defined by the sum of its diagonal elements, $\operatorname{tr}[X] = \sum_{j=1}^{\min(d_1,d_2)} X_{jj}$. It can be seen that the p-th power of the Schatten-p norm coincides with $\|X\|_{S_p}^p = \operatorname{tr}\left[(X^*X)^{p/2}\right]$. The Schatten-2 norm is also called *Frobenius norm* and has the property that it is induced by the Frobenius scalar product $\langle X, Y\rangle_F = \operatorname{tr}[X^*Y]$, i.e., $\|X\|_F = \|X\|_{S_2} = \sqrt{\langle X, X\rangle_F}$. We define the *weighted Frobenius scalar product* of two matrices $X, Y \in M_{d_1 \times d_2}$ weighted by the the

positive definite weight matrix $W \in H_{d_1 \times d_1}$ as $\langle X, Y \rangle_{F(W)} := \langle WX, Y \rangle_F = \langle X, WY \rangle_F$. This scalar product induces the *weighted Frobenius norm* $\|X\|_{F(W)} = \sqrt{\langle X, X \rangle_{F(W)}} = \sqrt{\mathrm{tr}[(WX)^* X]}$. A useful connection that will be used in the course of the paper is the equality of the Frobenius norm of a matrix X and the ℓ_2 norm of its vectorization X_{vec}, i.e., $\|X\|_F = \|X_{\mathrm{vec}}\|_{\ell_2}$.

Similar to weighted Frobenius norms, we define the *weighted ℓ_2-scalar product* of vectors $x, y \in \mathbb{C}^d$ weighted by the positive definite weight matrix $W \in H_{d \times d}$ as $\langle x, y \rangle_{\ell_2(W)} = x^* W y = y^* W x$ and its induced *weighted ℓ_2-norm* as $\|x\|_{\ell_2(W)} = \sqrt{x^* W x}$. We use the notation $X \succ 0$ for a positive definite matrix $X \in H_{d \times d}$. Furthermore, we denote the range of a linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ by $\mathrm{Ran}(\Phi) = \big\{ Y \in \mathbb{C}^m ;$ there is $X \in M_{d_1 \times d_2}$ such that $Y = \Phi(X) \big\}$ and its null space by $\mathcal{N}(\Phi) = \big\{ X \in M_{d_1 \times d_2} ; \Phi(X) = 0 \big\}$.

2.2 Kronecker and Hadamard products

For two matrices $A = (a_{ij})_{i \in [d_1], j \in [d_3]} \in \mathbb{C}^{d_1 \times d_3}$ and $B \in \mathbb{C}^{d_2 \times d_4}$, we call the matrix representation of their tensor product with respect to the standard bases the *Kronecker product* $A \otimes B \in \mathbb{C}^{d_1 \cdot d_2 \times d_3 \cdot d_4}$. By its definition, $A \otimes B$ is a block matrix of $d_2 \times d_4$ blocks whose block of index $(i, j) \in [d_1] \times [d_3]$ is the matrix $a_{ij} B \in \mathbb{R}^{d_2 \times d_4}$. This implies, e.g., for $A \in \mathbb{C}^{d_1 \times d_3}$ with $d_1 = 2$ and $d_3 = 3$ that

$$A \otimes B = \left[\begin{array}{ccc} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{array} \right] \otimes B = \left[\begin{array}{c|c|c} a_{11} B & a_{12} B & a_{13} B \\ \hline a_{21} B & a_{22} B & a_{23} B \end{array} \right].$$

For the reader's convenience we list some basic properties of the Kronecker product [3, Chapter 7], [53]:

1. $(A \otimes B)^* = A^* \otimes B^*$,
2. $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$ (whenever A and B are invertible),
3. $(A \otimes B)(C \otimes D) = (AC \otimes BD)$,
4. $A \otimes (B \otimes C) = (A \otimes B) \otimes C$.

We define the *Hadamard product* $A \circ B \in \mathbb{C}^{d_1 \times d_2}$ of two matrices $A \in \mathbb{C}^{d_1 \times d_2}$ and $B \in \mathbb{C}^{d_1 \times d_2}$ as their entry-wise product

$$(A \circ B)_{i,j} = A_{i,j} B_{i,j}$$

with $i \in [d_1]$ and $j \in [d_2]$. The Hadamard product is also known as *Schur product* in the literature. Again we list some of its basic calculation rules and refer to the literature for further properties [3, Chapter 7],

1. $A \circ B = B \circ A$,
2. $(A \circ B)^* = A^* \circ B^*$,
3. $A \circ (B + C) = A \circ B + A \circ C$,
4. $A \circ (B \circ C) = (A \circ B) \circ C$.

Furthermore, if $d_1 = d_3$ and $d_2 = d_4$, we define the *Kronecker sum* $A \oplus B \in \mathbb{C}^{d_1 d_2 \times d_1 d_2}$ of two matrices $A \in \mathbb{C}^{d_1 \times d_1}$ and $B \in \mathbb{C}^{d_2 \times d_2}$ as the matrix

$$A \oplus B = (\mathbf{I}_{d_2} \otimes A) + (B \otimes \mathbf{I}_{d_1}).$$

The Kronecker product is useful for the elegant formulation of matrix equations involving left and right matrix multiplications with the variable X, as

$$AXB^* = Y \qquad \text{if and only if} \qquad (B \otimes A) X_{\mathrm{vec}} = Y_{\mathrm{vec}}.$$

Kronecker sums can be used as a tool for the analysis of the Sylvester matrix equation problem, where one wants to find a matrix $X \in M_{d_1 \times d_2}$ that solves the equation

$$AX + XB^* = Y \tag{2}$$

for given $A \in \mathbb{C}^{d_1 \times d_1}, B \in \mathbb{C}^{d_2 \times d_2}$ and $Y \in \mathbb{C}^{d_1 \times d_2}$. This equation can be rewritten as

$$(A \oplus B)X_{\mathrm{vec}} = Y_{\mathrm{vec}},$$

using again the vectorizations of X and Y. In this context, we use an explicit formula that expresses the inverse $(A \oplus B)^{-1}$ of the Kronecker sum $A \oplus B$ (which should not be confused with the Kronecker *product* $A \otimes B$) in terms of singular vectors and singular values of A and B. More precisely, we use the following lemma.

Lemma 1 ([32]). *Let $A \in H_{d_1 \times d_1}$ and $B \in H_{d_2 \times d_2}$, where one of the matrices is positive definite and the other positive semidefinite. If we denote the singular vectors of A by $u_i \in \mathbb{C}^{d_1}$, $i \in [d_1]$, its singular values by σ_i, $i \in [d_1]$ and the singular vectors resp. values of B by $v_j \in \mathbb{C}^{d_2}$ resp. μ_j, $j \in [d_2]$, then*

$$(A \oplus B)^{-1} = \sum_{i=1}^{d_1} \sum_{j=1}^{d_2} \frac{v_j v_j^* \otimes u_i u_i^*}{\sigma_i + \mu_j} = (V \otimes U) D (V \otimes U)^*, \tag{3}$$

where $D \in M_{d_1 d_2 \times d_1 d_2}$ is a diagonal matrix with entries $d_l = (\sigma_i + \mu_j)^{-1} > 0$ for $l = (i-1)d_2 + j$, $U = [u_1, \ldots, u_{d_1}]$, and $V = [v_1, \ldots, v_{d_2}]$.

Furthermore, the action of $(A \oplus B)^{-1}$ on the matrix space $M_{d_1 \times d_2}$ can be written as

$$\big[(A \oplus B)^{-1} Z_{\mathrm{vec}}\big]_{\mathrm{mat}} = U\big(H \circ (U^* Z V)\big)V^*. \tag{4}$$

for $Z \in M_{d_1 \times d_2}$ and the matrix $H \in M_{d_1 \times d_2}$ with the entries $H_{i,j} = (\sigma_i + \mu_j)^{-1}$, $i \in [d_1]$, $j \in [d_2]$.

Lemma 1 can be shown by using the Kronecker product calculation rules indicated above.

2.3 Problem setting and characterization of S_p- and reweighted Frobenius norm minimizers

Given a linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ such that $m \ll d_1 d_2$, we want to uniquely identify and reconstruct an unknown matrix X_0 from its linear image $Y := \Phi(X_0) \in \mathbb{C}^m$. However, basic linear algebra tells us that this is not possible without further assumptions, since Φ is not injective if $m < d_1 d_2$. Indeed, there is a $(d_1 d_2 - m)$-dimensional affine space $\{X_0\} + \mathcal{N}(\Phi)$ fulfilling the linear constraint

$$\Phi(X) = Y.$$

Nevertheless, under the additional assumption that the matrix $X_0 \in M_{d_1 \times d_2}$ has rank $r < \min(d_1, d_2)$ and under appropriate assumptions on the map Φ, the recovery of X_0 is possible by solving the affine rank minimization problem

$$\min \operatorname{rank}(X) \text{ subject to } \Phi(X) = Y. \tag{5}$$

The unique solvability of (5) is given with high probability if, for example, Φ is a linear map whose matrix representation has i.i.d. Gaussian entries [15] and $m = \Omega(r(d_1 + d_2))$. Unfortunately, solving (5) is intractable in general, but the works [8], [43], and [7] suggest solving the tractable convex optimization program

$$\min \|X\|_{S_1} \text{ subject to } \Phi(X) = Y, \tag{6}$$

also called *nuclear norm minimization*, as a proxy.

As discussed in the introduction, there are empirical as well as theoretical results (e.g., in [48, 13, 10] coming from the related *sparse vector recovery* problem that suggest alternative relaxation approaches. These results indicate that it might be even more advantageous to solve the non-convex problem

$$\min F^p(X) := \|X\|_{S_p}^p \text{ subject to } \Phi(X) = Y, \tag{7}$$

for $0 < p < 1$, i.e., minimizing the p-th power of the Schatten-p quasi-norms under the affine constraint. Heuristically, the choice of $p < 1$ relatively small can be motivated by the observation that by the definition (1) of the Schatten-p quasinorm

$$\|X\|_{S_p}^p \xrightarrow{p\to 0} \operatorname{rank}(X) =: \|X\|_{S_0}.$$

The above consideration suggests that the solution of (7) might be closer to (5) rather than (6) for small p. On the other hand, again, it is in general computationally intractable to find a global minimum of the non-convex optimization problem (7) if $p < 1$. Therefore it is a natural and very relevant question to ask which optimization algorithm to use to find global minimizers of (7).

In this paper, we discuss an algorithm striving to solve (7) that is based on the following observations: Assume for the moment that we are given a square matrix $Z \in \mathbb{R}^{d_1\times d_2}$ with $d_1 = d_2$ of full rank. Then, we can rewrite the p-th power of its Schatten-p quasi-norm as squares of weighted Frobenius norms resp. weighted ℓ_2-norms (if we use the vectorized notation X_{vec}),

(i) $$\begin{aligned}\|Z\|_{S_p}^p &= \operatorname{tr}[(ZZ^*)^{\frac{p}{2}}] = \operatorname{tr}[(ZZ^*)^{\frac{p-2}{2}}(ZZ^*)] = \operatorname{tr}(W_L ZZ^*) = \|W_L^{\frac12} Z\|_F^2 = \|Z\|_{F(W_L)}^2 \\ &= \|(\mathbf{I}_{d_2}\otimes W_L)^{\frac12} Z_{\mathrm{vec}}\|_{\ell_2}^2 = \|Z_{\mathrm{vec}}\|_{\ell_2(\mathbf{I}_{d_2}\otimes W_L)}^2,\end{aligned}$$
where W_L is the symmetric weight matrix $(ZZ^*)^{\frac{p-2}{2}}$ in $M_{d_1\times d_1}$ and $\mathbf{I}_{d_2}\otimes W_L$ is the block diagonal weight matrix in $M_{d_1\cdot d_2\times d_1\cdot d_2}$ with d_2 instances of W_L on the diagonal blocks,

(ii) $$\begin{aligned}\|Z\|_{S_p}^p &= \operatorname{tr}[(Z^*Z)^{\frac{p}{2}}] = \operatorname{tr}[(Z^*Z)(Z^*Z)^{\frac{p-2}{2}}] = \operatorname{tr}(Z^*ZW_R) = \|ZW_R^{\frac12}\|_F^2 = \|Z^*\|_{F(W_R)}^2 \\ &= \|(W_R\otimes \mathbf{I}_{d_1})^{\frac12} Z_{\mathrm{vec}}\|_{\ell_2}^2 = \|Z_{\mathrm{vec}}\|_{\ell_2(W_R\otimes\mathbf{I}_{d_1})}^2,\end{aligned}$$
where W_R is the symmetric weight matrix $(Z^*Z)^{\frac{p-2}{2}}$ in $\mathbb{C}^{d_2\times d_2}$. The weight matrix $W_R\otimes\mathbf{I}_{d_1}\in\mathbb{C}^{d_1d_2\times d_1d_2}$ is, as can be calculated from the definition of the Kronecker product, a block matrix of diagonal blocks of the type $\operatorname{diag}((W_R)_{ij},\dots,(W_R)_{ij})\in\mathbb{C}^{d_1\times d_1}$, $i,j\in[d_2]$.

From (i) and (ii) it becomes clear that the p-th powers of the Schatten-p norm of Z can be expressed in terms of weighted Frobenius norms, by multiplying Z by $W_L^{\frac12}$ from the left or by $W_R^{\frac12}$ from the right, respectively.

The above calculations are not well-defined if Z is not of full rank or if $d_1 \neq d_2$, since in these cases at least one of the matrices $ZZ^*\in\mathbb{C}^{d_1\times d_1}$ or $Z^*Z\in\mathbb{C}^{d_2\times d_2}$ is singular, prohibiting the definition of the matrices $W_R = (Z^*Z)^{\frac{p-2}{2}}$ or $W_L = (ZZ^*)^{\frac{p-2}{2}}$ for $p < 2$. However, these issues can be overcome by introducing a smoothing parameter $\epsilon > 0$ and smoothed weight matrices $W_L(Z,\epsilon)\in\mathbb{C}^{d_1\times d_1}$ and $W_R(Z,\epsilon)\in\mathbb{C}^{d_2\times d_2}$ defined by

$$W_L(Z,\epsilon) := (ZZ^* + \epsilon^2\mathbf{I}_{d_1})^{\frac{p-2}{2}}, \tag{8}$$

$$W_R(Z,\epsilon) := (Z^*Z + \epsilon^2\mathbf{I}_{d_2})^{\frac{p-2}{2}}. \tag{9}$$

Remark 1. *The weight matrices $W_L(Z,\epsilon)$ and $W_R(Z,\epsilon)$ are symmetric and positive definite.*

2.4 Averaging of weight matrices

The left-sided reweighted Frobenius-norm $\|\cdot\|_{F(W_L)}^2$ reformulation of the Schatten-p norm was already considered for the deduction of iteratively reweighted least squares algorithms for low-rank matrix

recovery [39, 19]. We stress that both algorithms presented in these papers do not take advantage of the fact that the Schatten-p norm can also be reformulated using the reweighted Frobenius-norm with the weight matrix W_R multiplied from the right.

In this paper, we aim at combining both reweighting strategies into one to exploit the low rank information in the column as well as in the row space.

A first intuitive approach towards a symmetric exploitation of the low rank structure is inspired by the following identity.

Lemma 2. *Let $0 < p \leqslant 2$ and $Z \in M_{d_1 \times d_2}$ with $d = d_1 = d_2$ be a full rank matrix. Then*

$$\|Z\|_{S_p}^p = \frac{1}{2}\left(\|W_L^{\frac{1}{2}} Z\|_F^2 + \|Z W_R^{\frac{1}{2}}\|_F^2\right) = \left\| \left(\frac{W_L \oplus W_R}{2}\right)^{\frac{1}{2}} Z_{\mathrm{vec}} \right\|_{\ell_2}^2 = \|Z_{\mathrm{vec}}\|_{\ell_2(W_{(arith)})}^2,$$

where

$$\frac{1}{2}\left(\mathbf{I}_{d_2} \otimes W_L + W_R \otimes \mathbf{I}_{d_1}\right) = \frac{W_L \oplus W_R}{2} =: W_{(arith)}$$

is the arithmetic mean matrix of the symmetric and positive definite weight matrices $\mathbf{I}_{d_2} \otimes W_L$ and $W_R \otimes \mathbf{I}_{d_1}$, $W_L := (ZZ^*)^{\frac{p-2}{2}}$, and $W_R := (Z^*Z)^{\frac{p-2}{2}}$.

Proof. Using the computations above as well as the cyclicity of the trace, we calculate that

$$\begin{aligned}
\|Z\|_{S_p}^p &= \frac{1}{2}\left(\|W_L^{\frac{1}{2}} Z\|_F^2 + \|Z W_R^{\frac{1}{2}}\|_F^2\right) = \frac{1}{2}\left(\|(\mathbf{I}_{d_2} \otimes W_L)^{\frac{1}{2}} Z_{\mathrm{vec}}\|_{\ell_2}^2 + \|(W_R \otimes \mathbf{I}_{d_1})^{\frac{1}{2}} Z_{\mathrm{vec}}\|_{\ell_2}^2\right) \\
&= \frac{1}{2}\left[\mathrm{tr}\left((\mathbf{I}_{d_2} \otimes W_L)^{\frac{1}{2}} Z_{\mathrm{vec}} Z_{\mathrm{vec}}^* (\mathbf{I}_{d_2} \otimes W_L)^{\frac{1}{2}}\right) + \mathrm{tr}\left((W_R \otimes \mathbf{I}_{d_1})^{\frac{1}{2}} Z_{\mathrm{vec}} Z_{\mathrm{vec}}^* (W_R \otimes \mathbf{I}_{d_1})^{\frac{1}{2}}\right)\right] \\
&= \frac{1}{2}\left[\mathrm{tr}\left((\mathbf{I}_{d_2} \otimes W_L) Z_{\mathrm{vec}} Z_{\mathrm{vec}}^*\right) + \mathrm{tr}\left((W_R \otimes \mathbf{I}_{d_1}) Z_{\mathrm{vec}} Z_{\mathrm{vec}}^*\right)\right] \\
&= \frac{1}{2}\,\mathrm{tr}\left(\left[(\mathbf{I}_{d_2} \otimes W_L) + (W_R \otimes \mathbf{I}_{d_1})\right] Z_{\mathrm{vec}} Z_{\mathrm{vec}}^*\right) = \frac{1}{2}\left\|(\mathbf{I}_{d_2} \otimes W_L + W_R \otimes \mathbf{I}_{d_1})^{\frac{1}{2}} Z_{\mathrm{vec}}\right\|_{\ell_2}^2 \\
&= \left\| \left(\frac{W_L \oplus W_R}{2}\right)^{\frac{1}{2}} Z_{\mathrm{vec}} \right\|_{\ell_2}^2.
\end{aligned}$$

□

Unfortunately, the introduction of arithmetic mean weight matrices does not prove to be particularly advantageous compared to one-sided reweighting strategies. No convincing improvements can be noted neither in numerical experiments nor in the theoretical investigations for the convergence rate of IRLS for low-rank matrix recovery, cf. also subsection 5.2 and Remark 7.

In contrast, the usage of the harmonic mean of the weight matrices $\mathbf{I}_{d_2} \otimes W_L$ and $W_R \otimes \mathbf{I}_{d_1}$, i.e., weight matrices of the type $2\left(W_R^{-1} \otimes \mathbf{I}_{d_1} + \mathbf{I}_{d_2} \otimes W_L^{-1}\right)^{-1} = 2\left(W_L^{-1} \oplus W_R^{-1}\right)^{-1} =: W_{(\mathrm{harm})}$ turn out to significantly outperform other weighting variants both theoretically and practically as shown in the upcoming sections.

The following lemma verifies that also the harmonic mean of the weight matrices $\mathbf{I}_{d_2} \otimes W_L$ and $W_R \otimes \mathbf{I}_{d_1}$ leads to a legitimate reformulation of the Schatten-p quasi-norm power.

Lemma 3. *Let $0 < p \leqslant 2$ and $Z \in \mathbb{C}^{d_1 \times d_2}$ with $d = d_1 = d_2$ be a full rank matrix. Then*

$$\|Z\|_{S_p}^p = 2\left\|\left(W_L^{-1} \oplus W_R^{-1}\right)^{-\frac{1}{2}} Z_{\mathrm{vec}}\right\|_{\ell_2}^2 = \|Z_{\mathrm{vec}}\|_{\ell_2(W_{(harm)})}^2,$$

where

$$2\left(W_R^{-1} \otimes \mathbf{I}_{d_1} + \mathbf{I}_{d_2} \otimes W_L^{-1}\right)^{-1} = 2\left(W_L^{-1} \oplus W_R^{-1}\right)^{-1} =: W_{(harm)}$$

is the harmonic mean matrix of the symmetric and positive definite weight matrices $\mathbf{I}_{d_2} \otimes W_L$ and $W_R \otimes \mathbf{I}_{d_2}$, $W_L := (ZZ^*)^{\frac{p-2}{2}}$ and $W_R := (Z^*Z)^{\frac{p-2}{2}}$.

Proof. Let $Z = USV^* \in M_{d\times d}$ be the singular value decomposition of Z. Therefore for the vectorized version $Z_{\mathrm{vec}} = (V \otimes U)S_{\mathrm{vec}}$ holds true. By the definitions of W_L and W_R, we can write $W_L^{-1} = US^{2-p}U^*$ and $W_R^{-1} = VS^{2-p}V^*$. Using the Kronecker sum inversion formula (3), we obtain $\left(W_L^{-1} \oplus W_R^{-1}\right)^{-1} = (V \otimes U)D(V \otimes U)^*$, where $D \in M_{d_1d_2\times d_1d_2}$ is a diagonal matrix with entries $d_l = (s_i^{2-p} + s_j^{2-p})^{-1} > 0$ for $l = (i-1)d_2 + j$, if $i \in [d_1]$ and $j \in [d_2]$.

Having these facts at hand, we calculate using the orthonormality of the columns of U and of V and the special structure of D, that

$$\begin{aligned}
\|Z_{\mathrm{vec}}\|^2_{\ell_2(W_{(\mathrm{harm})})} &= \|W_{(\mathrm{harm})}^{\frac{1}{2}} Z_{\mathrm{vec}}\|^2_{\ell_2} = 2\left\|\left(W_L^{-1} \oplus W_R^{-1}\right)^{-\frac{1}{2}} Z_{\mathrm{vec}}\right\|^2_{\ell_2} \\
&= 2\,\mathrm{tr}\left(\left(\left(W_L^{-1} \oplus W_R^{-1}\right)^{-1} Z_{\mathrm{vec}}\right)^*_{\mathrm{mat}} Z\right) \\
&= 2\,\mathrm{tr}\left(\left[(V \otimes U)D(V \otimes U)^*(V \otimes U)S_{\mathrm{vec}}\right]^* (V \otimes U)S_{\mathrm{vec}}\right) \\
&= 2\,\mathrm{tr}\left(S^*_{\mathrm{vec}} D S_{\mathrm{vec}}\right) \\
&= 2\left(\sum_{i=1}^{d} \frac{s_i^2}{2s_i^{2-p}}\right) = \|Z\|^p_{S_p}.
\end{aligned}$$

□

3 Harmonic mean iteratively reweighted least squares algorithm

In this section, we formulate a new iteratively reweighted least squares algorithm for low-rank matrix recovery. The so-called *harmonic mean iteratively reweighted least squares algorithm* (`HM-IRLS`) solves a sequence of weighted least squares problems to recover a low-rank matrix $X_0 \in M_{d_1\times d_2}$ from few linear measurements $\Phi(X_0) \in \mathbb{C}^m$. The weight matrices appearing in the least squares problems can be seen as the harmonic mean of the weight matrices in (8) and (9).

Let $0 < p \leqslant 1$ and denote $d = \min(d_1, d_2), D = \max(d_1, d_2)$. We propose the following approach: Given a non-increasing sequence of non-negative real numbers $(\epsilon^{(n)})_{n=1}^{\infty}$, we initialize a (positive definite) weight matrix $\widetilde{W}^{(0)} \in H_{d_1d_2\times d_1d_2}$. Then we recursively define for $n = 1, 2, \ldots,$

$$X^{(n)} = \operatorname*{arg\,min}_{\Phi(X)=Y} \|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}^{(n-1)})} \tag{10}$$

and

$$\widetilde{W}^{(n)} = 2\left[U^{(n)}(\bar{\Sigma}^{(n)}_{d_1})^{2-p}U^{(n)*} \oplus V^{(n)}(\bar{\Sigma}^{(n)}_{d_2})^{2-p}V^{(n)*}\right]^{-1} \tag{11}$$

with the diagonal matrices $\bar{\Sigma}^{(n)}_{d_t} \in M_{d_t\times d_t}$ for $d_t = \{d_1, d_2\}$ such that

$$(\bar{\Sigma}^{(n)}_{d_t})_{ii} = \begin{cases} (\sigma_i(X^{(n)})^2 + \epsilon^{(n)2})^{\frac{1}{2}} & \text{if } i \leqslant d, \\ 0 & \text{if } d < i \leqslant D \end{cases} \tag{12}$$

and the matrices $U^{(n)} \in M_{d_1\times d_1}$ and $V^{(n)} \in M_{d_2\times d_2}$, containing the left and right singular vectors of $X^{(n)}$ in its columns, respectively.

We note that this definition of $\widetilde{W}^{(n)}$ can be seen as a stabilized version of the harmonic mean weight matrix $W_{(\mathrm{harm})}$ of Lemma 3. This stabilization is necessary as $\widetilde{W}^{(n)}$ becomes very ill-conditioned as soon as some of the singular values of $X^{(n)}$ approach zero and, related to that, $(X^{(n)}X^{(n)*})^{\frac{2-p}{2}} \oplus (X^{(n)*}X^{(n)})^{\frac{2-p}{2}}$ would be even singular as soon as $X^{(n)}$ is not of full rank.

To write down the first step (10) of HM-IRLS in a concise way, we define for $n \in \mathbb{N}$ the linear operator $\widetilde{\mathcal{W}}^{(n)-1} : M_{d_1 \times d_2} \to M_{d_1 \times d_2}$ as

$$(\widetilde{\mathcal{W}}^{(n)-1})(Z) := \frac{1}{2}\left[U^{(n)}(\bar{\Sigma}_{d_1}^{(n)})^{2-p}U^{(n)*}Z + ZV^{(n)}(\bar{\Sigma}_{d_2}^{(n)})^{2-p}V^{(n)*}\right], \tag{13}$$

which describes the operation of the inverse of $\widetilde{W}^{(n)}$ on $M_{d_1 \times d_2}$.

In pseudo code, the approach of HM-IRLS can be summarized as follows.

Algorithm 1: Harmonic Mean IRLS for low-rank matrix recovery (HM-IRLS)

Input: A linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$, image $Y = \Phi(X_0)$ of the ground truth matrix $X_0 \in M_{d_1 \times d_2}$, rank parameter $\widetilde{r}$, non-convexity parameter $0 < p \leqslant 1$.

Output: Sequence $(X^{(n)})_{n=1}^{n_0} \subset M_{d_1 \times d_2}$.

Initialize $n = 0$, $\epsilon^{(0)} = 1$ and $\widetilde{W}^{(0)} = \mathbf{I}_{d_1 d_2} \in M_{d_1 d_2 \times d_1 d_2}$.

repeat

$$X^{(n+1)} = \operatorname*{arg\,min}_{\Phi(X)=Y} \|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}^{(n)})} = \big(\widetilde{\mathcal{W}}^{(n)-1} \circ \Phi^* \circ (\Phi \circ \widetilde{\mathcal{W}}^{(n)-1} \circ \Phi^*)^{-1}\big)\big(Y\big), \tag{14}$$

$$\epsilon^{(n+1)} = \min\left(\epsilon^{(n)}, \sigma_{\widetilde{r}+1}(X^{(n+1)})\right), \tag{15}$$

$$\widetilde{W}^{(n+1)} = 2\left[U^{(n+1)}(\bar{\Sigma}_{d_1}^{(n+1)})^{2-p}U^{(n+1)*} \oplus V^{(n+1)}(\bar{\Sigma}_{d_2}^{(n+1)})^{2-p}V^{(n+1)*}\right]^{-1}, \tag{16}$$

where $U^{(n+1)} \in M_{d_1 \times d_1}$ and $V^{(n+1)} \in M_{d_2 \times d_2}$ are matrices containing the left and right singular vectors of $X^{(n+1)}$ in its columns, and $\bar{\Sigma}^{(n+1)}$ is defined according to (12).

$n = n + 1$.

until *stopping criterion is met.*

Set $n_0 = n$.

From a practical point of view, it is beneficial that the explicit calculation of the very large weight matrices $\widetilde{W}^{(n+1)} \in H_{d_1 d_2 \times d_1 d_2}$ (cf. (16)) is not necessary in implementations of Algorithm 1. As suggested by formulas (13) and (14), it can be seen that just the operation of *its inverse* $(\widetilde{W}^{(n+1)})^{-1}$ *resp.* $(\widetilde{W}^{(n)})^{-1}$ is needed, which can be implemented by matrix-matrix multiplications on the space $M_{d_1 \times d_2}$: For matrices $X, Z \in M_{d_1 \times d_2}$, we have that $\widetilde{W}^{(n)}X_{\mathrm{vec}} = Z_{\mathrm{vec}}$ if and only if $X_{\mathrm{vec}} = (\widetilde{W}^{(n)})^{-1}Z_{\mathrm{vec}}$, which can be written in matrix variables as

$$X = \frac{1}{2}\left[U^{(n)}(\bar{\Sigma}_{d_1}^{(n)})^{2-p}U^{(n)*}Z + ZV^{(n)}(\bar{\Sigma}_{d_2}^{(n)})^{2-p}V^{(n)*}\right].$$

The last equivalence is due to the definition of $\widetilde{W}^{(n)}$ (cf. (16)) and the property (2) of the Kronecker sum.

We refer to section 5.4 for a further discussion of implementation details.

4 Convergence results

In the following part, we state our main theoretical results about convergence properties of the algorithm HM-IRLS. Furthermore, their relation to results for existing IRLS-type algorithms is discussed.

4.1 Local convergence for $p < 1$

To keep our exposition accessible, we first consider the special case of a random linear measurement operator $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ whose matrix representation is

$$\frac{1}{\sqrt{m}}\widetilde{\Phi} \in \mathbb{C}^{m \times d_1 d_2}, \text{ where } \widetilde{\Phi} \text{ has i.i.d. standard Gaussian entries.} \tag{17}$$

In this case, we call Φ a *Gaussian operator*.

Next, for a given matrix $X \in M_{d_1 \times d_2}$, $0 < p \leqslant 1$ and $\epsilon > 0$, we define the *ϵ-perturbed Schatten-p functional* $g_\epsilon^p : M_{d_1 \times d_2} \to \mathbb{R}_{\geqslant 0}$ such that

$$g_\epsilon^p(X) = \sum_{i=1}^{d} (\sigma_i(X)^2 + \epsilon^2)^{\frac{p}{2}} \tag{18}$$

where $\sigma(X) \in \mathbb{R}^d$ denotes the vector of singular values of X.

Our convergence result for `HM-IRLS` for the special case of Gaussian operators can now be formulated as follows:

Theorem 1. *Let $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ be a Gaussian operator and let $Y \in \operatorname{Ran}(\Phi)$ be a vector in its range. Let $(X^{(n)})_{n \geqslant 1}$ and $(\epsilon^{(n)})_{n \geqslant 1}$ be the sequences produced by Algorithm 1 for input parameters Φ, Y, r and $0 < p \leqslant 1$, let $\epsilon = \lim_{n \to \infty} \epsilon^{(n)}$.*

(i) There exists a constant $\rho_1 \geqslant 1$ such that if $m > \rho_1 r(d_1 + d_2)$, with high probability, the following holds:

If $\epsilon = 0$, then the sequence $(X^{(n)})_{n \geqslant 1}$ converges to a matrix $\bar{X} \in M_{d_1 \times d_2}$ of rank at most r that is the unique minimizer of the Schatten-p minimization problem (7). Moreover, there exists an absolute constant $\hat{C} > 0$ such that for any Z with $\Phi(Z) = Y$ and any $\widetilde{r} \leqslant r$, it holds that

$$\|Z - \bar{X}\|_F^p \leqslant \frac{\hat{C}}{r^{1-\frac{p}{2}}} \beta_{\widetilde{r}}(Z)_{S_p}, \tag{19}$$

where $\beta_{\widetilde{r}}(Z)_{S_p}$ is the best rank-$\widetilde{r}$ Schatten-p approximation error of Z, i.e.,

$$\beta_{\widetilde{r}}(Z)_{S_p} := \inf\left\{\|Z - \widetilde{Z}\|_{S_p}^p,\ \widetilde{Z} \in M_{d_1 \times d_2} \text{ has rank } \widetilde{r}\right\}.$$

(ii) If $\epsilon > 0$, then each accumulation point $\bar{X}$ of $(X^{(n)})_{n \geqslant 1}$ is a stationary point of the ϵ-perturbed Schatten-p functional g_ϵ^p of (18) under the linear constraint $\Phi(X) = Y$. If additionally $p = 1$, then $\bar{X}$ is the unique global minimizer of g_ϵ^p.

(iii) There exists a constant $\rho_2 \geqslant 1$ such that if $m > \rho_2 r(d_1 + d_2)$, with high probability, the following holds:

Assuming that there exists a matrix $X_0 \in M_{d_1 \times d_2}$ with $\Phi(X_0) = Y$ of rank $\widetilde{r} \leqslant r$, a constant $0 < \rho < 1$ and an iteration $\bar{n} \in \mathbb{N}$ such that

$$\|X^{(\bar{n})} - X_0\|_{S_\infty} \leqslant \rho \sigma_{\widetilde{r}}(X_0) \tag{20}$$

and $\epsilon^{\bar{n}} = \sigma_{r+1}(X^{\bar{n}})$.

If the condition number $\kappa = \frac{\sigma_1(X_0)}{\sigma_{\widetilde{r}}(X_0)}$ of X_0 and ρ are sufficiently small (see condition (22) and formula (69)), then

$$X^{(n)} \to X_0 \quad \text{for } n \to \infty.$$

The proof and a more general version of Theorem 1 can be found in section 6.4, cf. Theorem 4. Thus, it can be said that the conclusions of Theorem 1 apply to more general measurement operators Φ than Gaussians, more specifically, to measurement operators that fulfill certain Schatten-p null space properties. These null space properties are fulfilled with high probability for a wide class of random operators in the optimal sample complexity regime, e.g. for operators with i.i.d. subgaussian entries.

4.2 Locally superlinear convergence rate for $p < 1$

Next, we state the second main theoretical result of this paper, Theorem 2. It shows that in a neighborhood of a low-rank matrix X_0 that is compatible with the measurement vector Y, the algorithm `HM-IRLS` converges to X_0 with a*convergence rate* that is *superlinear of the order* $2 - p$. Similarly to Theorem 1, we state it for Gaussian operators for presentation reasons, but note that it generalizes to measurement operators fulfilling the Schatten-p null space property, cf. Theorem 5. The theorems are proven in Section 6.5.

Theorem 2. *Let $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ be a Gaussian operator and $Y \in \mathrm{Ran}\,(\Phi)$ be a vector in its range. Assume that there exists a matrix $X_0 \in M_{d_1 \times d_2}$ with $\mathrm{rank}(X_0) = r \leqslant \frac{\min(d_1,d_2)}{2}$ such that $\Phi(X_0) = Y$, and let $(X^{(n)})_{n \geqslant 1}$ be the output sequence of Algorithm 1 with input parameters Φ, Y, r and $0 < p \leqslant 1$. Moreover, let $\kappa = \frac{\sigma_1(X_0)}{\sigma_r(X_0)}$ be the condition number of X_0 and $\eta^{(n)} := X^{(n)} - X_0$ be the residuals of the n-th output matrix of Algorithm 1 for $n \in \mathbb{N}$.*

Then, there exists a constant $\rho_2 \geqslant 1$ such that if $m > \rho_2 r(d_1 + d_2)$, the following holds with high probability:

Given an iteration $\bar{n} \in \mathbb{N}$ and a constant $0 < \rho < 1$ such that

$$\|\eta^{(\bar{n})}\|_{S_\infty} \leqslant \rho \sigma_r(X_0) \tag{21}$$

and $\epsilon^{\bar{n}} = \sigma_{r+1}(X^{\bar{n}})$.
If additionally the condition number κ is small enough such that the constant μ as specified in formula (69) *fulfills*

$$\mu \|\eta^{(\bar{n})}\|_{S_\infty}^{p(1-p)} < 1, \tag{22}$$

then

$$\|\eta^{(n+1)}\|_{S_\infty} \leqslant \mu^{1/p} \left(\|\eta^{(n)}\|_{S_\infty}\right)^{2-p} \quad \text{and} \quad \|\eta^{(n+1)}\|_{S_p} \leqslant \mu^{1/p} \left(\|\eta^{(n)}\|_{S_p}\right)^{2-p}$$

for all $n \geqslant \bar{n}$.

Remark 2. *It is interesting to compare Theorem 2 with a related result for an IRLS algorithm for the sparse vector recovery problem [13, Theorem 7.9]. We observe that while the statement describes the observed rates of convergence very accurately (cf. subsection 5.2), the assumption* (22) *on the neighborhood that enables convergence of a rate $2 - p$ is more pessimistic than our numerical experiments suggest.*

At this point, it is important to mention that although the convergence guarantees of our theory apply to wide class of more general measurement operators as proven in section 6, they currently do not comprise the matrix completion setting, i.e., measurement operators $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ such that

$$\Phi(X)_\ell = X_{i_\ell, j_\ell}, \quad \ell = 1, \ldots, m, \tag{23}$$

where $\Omega = \{(i_\ell, j_\ell); \ell \in [m]\} \subset [d_1] \times [d_2]$ is a fixed set of m index pairs. This is due to the fact that matrix completion measurement operators do not fulfill Schatten-p null space properties if $m < d_1 d_2$, cf. Remark 6 of section 6.4.

However, in section 5, we provide strong numerical evidence that the conclusions of our theorems are true even for the matrix completion setting.

4.3 Discussion and comparison with existing IRLS algorithms

Optimally, we would like to have a statement in Theorem 1 about the accumulation points $\bar{X}$ being *global minimizers* of g_ϵ^p, instead of mere stationary points, cf. [19, Theorem 6.11], [13, Theorem 5.3]. A statement that strong is, unfortunately, difficult to achieve due to the non-convexity of the Schatten-p quasinorm and of the ϵ-perturbed version g_ϵ^p. Nevertheless, our theorems can be seen as

analogues of [13, Theorem 7.7], which discusses the convergence properties of an IRLS algorithm for sparse recovery based on ℓ_p-minimization with $p < 1$.

As already mentioned in section 1, the authors of [19] and [39] proposed IRLS algorithms for low-rank matrix recovery and analysed their convergence properties. The algorithm of [19] can be seen as an asymmetric variant of `HM-IRLS` corresponding to the choice $p = 1$ in our notation. In this context, Theorem 1 recovers the results of [19, Theorem 6.11(i-ii)] for $p = 1$ and generalize them, with weaker conclusions due to the non-convexity, to the cases $0 < p < 1$. The algorithm `IRLS-`p of [39] is similar to the former, but also covers non-convex choices $0 < p < 1$. However, we note that in the non-convex case, its convergence result [39, Theorem 5.1] corresponds to Theorem 1(ii), but does not provide statements similar to (i) and (iii) of Theorem 1.

Theorem 2 with its analysis of the convergence rate is new in the sense that to the best of our knowledge, there are no convergence rate proofs for IRLS algorithms for the low-rank matrix recovery problem in the literature. Indeed, we refer to Remark 7 in section 6.5 for an explanation why the variants of [19] and [39] cannot exhibit superlinear convergence rates, unlike `HM-IRLS`.

We also note that there is a close connection between the statements of Theorems 1 and 2 and results that were obtained for an IRLS algorithm dedicated to the sparse vector recovery problem [13, Theorems 7.7 and 7.9].

5 Numerical experiments

In this section, we demonstrate that the superlinear convergence rate that was proven theoretically for Algorithm 1 (`HM-IRLS`) in Theorem 2 can indeed be accurately verified in numerical experiments even beyond measurement operators fulfilling the strong null space property. In particular, the experiments detailed in subsection 5.2 show that `HM-IRLS` exhibits a superlinear convergence rate of order $2 - p$ that can be observed very clearly in the important matrix completion setting. Additionally, we show that a superlinear rate of any kind can *not* be observed for other variants of IRLS [19, 39], including a variant that uses arithmetic mean weight matrices (see Lemma 2 for a motivation) instead of harmonic mean weight matrices.

In subsection 5.3, we examine the recovery performance of `HM-IRLS` for the matrix completion setting with the performance of other algorithms, comparing the measurement complexities that are needed for successful recovery for many random instances. The algorithms that are compared to `HM-IRLS` in the experiments include variants of the iteratively reweighted least squares algorithm as well as other types of state-of-the-art low-rank matrix recovery methods. Interestingly, the experiments show that despite the severe non-convexity of the underlying problem, `HM-IRLS` is able to recover the low-rank matrix even for cases when only a very low sample rate is available ($\rho \approx 1$), without the need of a special initialization. In particular, `HM-IRLS` *recovers low-rank matrices systematically with nearly the optimal number of measurements and needs fewer measurements than all the state-of-the-art algorithms, including previously existing IRLS methods, involved in our experiments.*

All numerical experiments discussed in this paper are performed on a MacBook Pro 9.1 with a 2.6 GHz Intel Core i7 quad-core-processor and 8GB memory. Computations are run in MATLAB R2014a, version 8.3.0.532. An implementation of the `HM-IRLS` algorithm and a minimal test example are available at https://www-m15.ma.tum.de/Allgemeines/SoftwareSite.

5.1 Measurement setting

We consider $d_1 \times d_2$ low-rank matrices X_0 of $\operatorname{rank}(X_0) = r$ that are constructed by the multiplication of matrices $U\Sigma V$, where $U \in \mathbb{R}^{d_1 \times r}$, $V \in \mathbb{R}^{r \times d_2}$ are matrices with independent standard Gaussian entries and $\Sigma \in R^{r \times r}$ is a diagonal matrix with i.i.d. standard Gaussian entries as well.

We note that a low-rank matrix $X \in M_{d_1 \times d_2}$ of rank r has $d_f = r(d_1 + d_2 - r)$ degrees of freedom, which is the theoretical lower bound on the number of measurements that are necessary for exact reconstruction [7]. The random measurement setting we use in the experiments can be described

as follows: We take measurements of matrix completion type, sampling $m = \lfloor \rho \cdot d_f \rfloor$ entries of X_0 uniformly over its $d_1 \cdot d_2$ indices to obtain $Y = \Phi(X_0)$. Here, the parameter $\frac{d_1 d_2}{d_f} \geqslant \rho \geqslant 1$ regulates the hardness of the reconstruction problem.

However, sampling Φ like that would yield instances of measurement operators whose information content is not large enough to ensure well-posedness of the corresponding low-rank matrix recovery problem, even if $\rho > 1$. More precisely, it is impossible to recover a matrix exactly if the number of revealed entries in any row or column is smaller than its rank r, which is explained and shown in the context of the proof of [42, Theorem 1].

Thus, in order to provide for a sensible measurement model for small ρ, we exclude operators Φ that sample fewer than r entries in any row or column. Therefore, we adapt the uniform sampling model such that operators Φ are discarded and generated again until the requirement of at least r entries per column/row is met and recovery can be achieved from a theoretical point of view.

We note that the described phenomenon is very related to the fact that matrix completion recovery guarantees for the uniform sampling model require at least one additional log factor, i.e., they require at least $m \geqslant \log(\max(d_1, d_2))d_f$ sampled entries [14, Section V].

While we detail the experiments for the matrix completion measurement setting just described in the remaining section, we add that Gaussian measurement models also lead to very similar results in experiments.

5.2 Convergence rate comparison with other IRLS type algorithms

In this subsection, we compare the harmonic mean iteratively reweighted least squares algorithm to previous variants of IRLS for matrix recovery as developed in [19] (IRLS-M) denoted by `IRLS-FRW`, and the very similar version [39] (IRLS-p) named `IRLS-MF`, which all only use reweightings in the column space. Moreover, we compare the performance of `HM-IRLS` with an alternative approach of reweightings in both the row and column space: We include an algorithm called arithmetic mean iteratively reweighted least squares (`AM-IRLS`), which is similar to Algorithm 1, but uses weight matrix composed according to the arithmetic mean

$$\widetilde{W}^{(n)} = \frac{1}{2}\left[\left(U^{(n)} S^{(n)} S^{(n)*} U^{(n)*} + \epsilon^{(n)2} \mathbf{I}_{d_1}\right)^{\frac{p-2}{2}} \oplus \left(V^{(n)} S^{(n)*} S^{(n)} V^{(n)*} + \epsilon^{(n)2} \mathbf{I}_{d_2}\right)^{\frac{p-2}{2}} \right]$$

of left- and right-sided weight matrices instead of (11), if $X^{(n)} = U^{(n)} S^{(n)} V^{(n)*}$ is the full SVD of an iterate $X^{(n)}$. For the motivation of `AM-IRLS`, see subsection 2.4.

The goal of the comparisons is the study of the convergence behavior of the four algorithms for hard and easy matrix completion problems, including their convergence speeds. For the experiments, we choose the parameters $d_1 = d_2 = 40$, $r = 10$, $\rho = 2.0, 1.2, 1.0$ and examine instances of the random model described in subsection 5.1. We conduct the experiments for non-convexity parameters $p = \{0.0001, 0.05, 0.1, 0.25, 0.5, 0.65, 0.8, 1.0\}$ related to the Schatten-p minimization problems (7) that all the four IRLS algorithms strive to solve.

Therefore, we plot the behavior of the logarithm of the Frobenius error $\log(\|X^{(n)} - X_0\|_F)$ for the iterations $n = 1, \ldots, \bar{n}$ of the mentioned algorithms in Figures 1 to 3, where $\bar{n}$ is the iteration when either a certain Frobenius error tolerance is undercut or a maximum number of iterations $n_{\max}$ is reached.

5.2.1 Results for `HM-IRLS`

In Figure 1a, we observe that `HM-IRLS` exhibits superlinear rates of convergence for choices of $p < 1$ exactly of the orders $2 - p$ that are predicted by Theorem 2. If the parameter p tends from 1 to 0, this entails a dramatic improvement of the convergence rate *from linear to (nearly) quadratic.* Increasing the hardness of the matrix completion problem by decreasing ρ from $\rho = 2$ in Figures 1 to $\rho = 1.2$ in Figure 2 and finally even to $\rho = 1$ in 3 leads to divergence of `HM-IRLS` for large values $p \gg 0$. This is

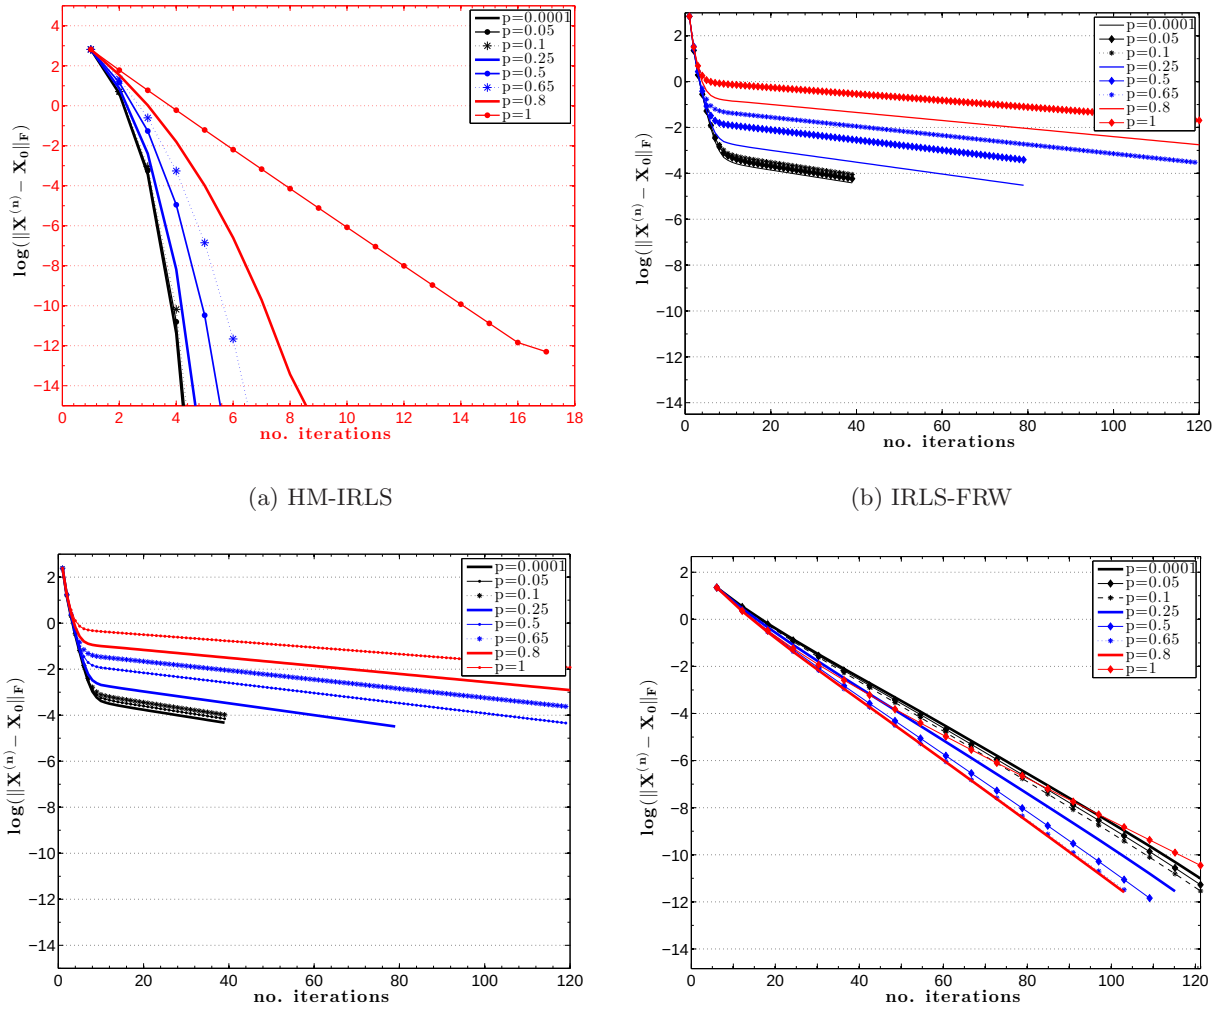

(a) HM-IRLS

(b) IRLS-FRW

(c) IRLS-MF

(d) AM-IRLS

Figure 1: Behavior of the log error $\log(\|X^{(n)} - X_0\|_F)$ for successive iterations for different parameter values of $0 < p \leqslant 1$ and fixed measurement oversampling factor $\rho = 2.0$. Note that the x-axis in (a) has a different scaling, indicating a much faster convergence of `HM-IRLS`.

not surprising at all, since a choice of $p = 1$ means that `HM-IRLS` just approximates the nuclear norm minimizer, which is known not to be able to recover the low-rank matrix if the sample complexity ρ is close to 1.

However, we observe that very interestingly, by choosing p close to 0 in the regimes of very low sample complexities $\rho = 1.2$ and $\rho = 1$, `HM-IRLS` retains the successful recovery of the low-rank matrix and still exhibits the convergence rate of order $2 - p$. This can be seen in Figures 2a and 3.

5.2.2 Results for other IRLS-type algorithms for low rank matrix recovery

In contrast, we observe that the variants `IRLS-FRW`, `IRLS-MF` and `AM-IRLS` *do not converge* to the ground truth low-rank matrix X_0 for hard matrix completion problems of low measurement complexity parameters $\rho = 1.2$ and $\rho = 1.0$, regardless of the choice of p. This can be seen in Figures 2b–2d, the

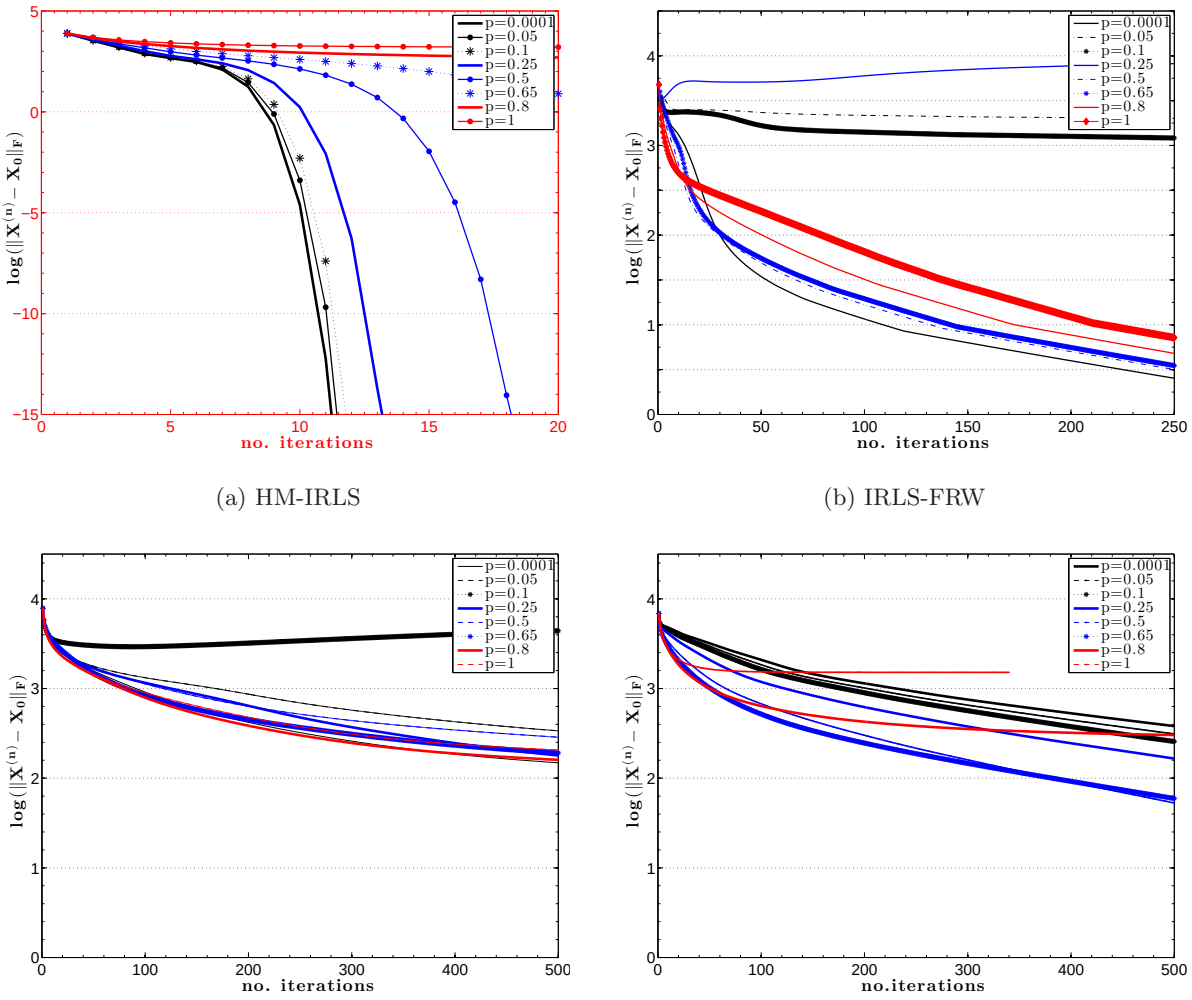

(a) HM-IRLS

(b) IRLS-FRW

(c) IRLS-MF

(d) AM-IRLS

Figure 2: Behavior of the log error $\log(\|X^{(n)} - X_0\|_F)$ for successive iterations for different parameter values of $0 < p \leqslant 1$ and fixed measurement oversampling factor $\rho = 1.2$. Note that the x-axis in (a) has a different scaling, indicating a much faster convergence of `HM-IRLS`.

corresponding plots for $\rho = 1.0$ were not included in Figure 3 due to a similar lack of convergence.

Whenever they converge, which is the case for easier problems corresponding to $\rho = 2.0$, they exhibit at best only a *linear* rate of convergence, even for small parameters $p \ll 1$, which can be seen in Figure 1b–1d. Regarding the non-convexity parameter p, it can be said that `IRLS-FRW` and `IRLS-MF` converge slightly faster for p approaching 0 if a generous amount of measurements is provided, while larger values of p give more promising results for `AM-IRLS` in this case (cf. again Figure 1b–1d).

For problems of intermediate difficulty corresponding to $\rho = 1.2$, we note that `IRLS-FRW`, `IRLS-MF` and `AM-IRLS` become unstable very easily if p is chosen too small. We conjecture that there exists an optimal value $0 < p_{\text{opt}} < 1$ for each algorithm, but did not investigate this further.

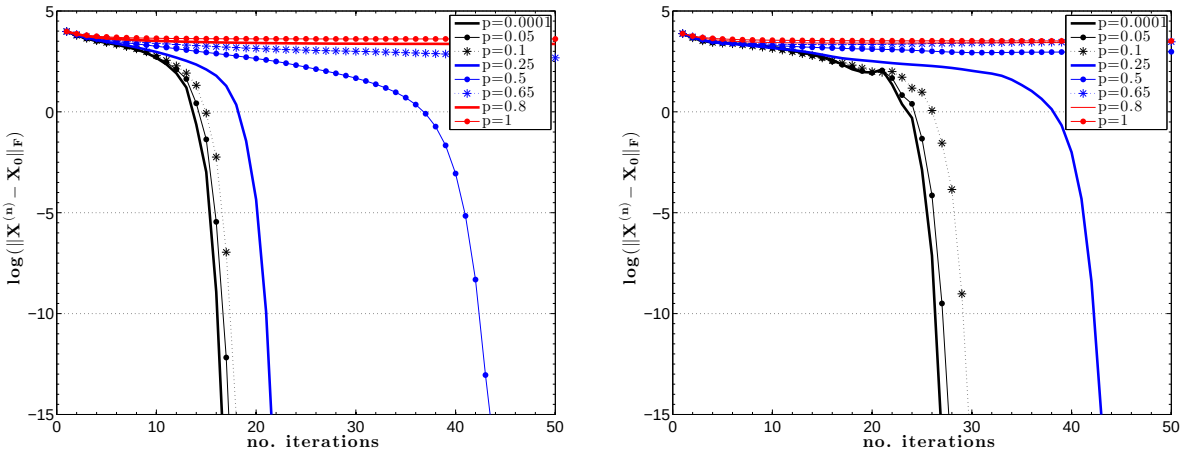

(a) HM-IRLS

(b) HM-IRLS

Figure 3: Behavior of the log error $\log(\|X^{(n)} - X_0\|_F)$ for successive iterations for different parameter values of $0 < p \leqslant 1$ for measurement oversampling factor $\rho = 1.0$, two different instances of the measurement model of section 5.1.

5.2.3 `HM-IRLS` as the best extension of IRLS for sparse recovery

We summarize that among the four variants `HM-IRLS`, `IRLS-FRW`, `IRLS-MF` and `AM-IRLS`, only `HM-IRLS` is able to solve the low-rank matrix recovery problem for very low sample complexities corresponding to $\rho \approx 1$. Furthermore, it is the only IRLS algorithm for low-rank matrix recovery that exhibits a superlinear rate of convergence at all.

It is worthwhile to compare the properties of `HM-IRLS` with the behavior of the IRLS algorithm of [13] designed to solve the sparse vector recovery problem by mimicking ℓ_p-minimization for $0 < p \leqslant 1$. While neither `IRLS-FRW`, `IRLS-MF` nor `AM-IRLS` are able to generalize the superlinear convergence behavior of [13] as illustrated in Figure 8.3 in [13] to the low-rank matrix recovery problem, `HM-IRLS` is, as can be seen in Figures 1a, 2a and 3.

Taking the theoretical guarantees as well as the numerical evidence into account, we claim that *`HM-IRLS` is the presently best extension of IRLS for vector recovery [13] to the low-rank matrix recovery setting*, providing a substantial improvement over the reweighting strategies of [19, 39].

Moreover, we mention two observations which suggest that `HM-IRLS` has in some sense even more favorable properties than the algorithm of [13]: First, the discussion of [13, Section 8] states that a superlinear convergence can only be observed locally after a considerable amount of iterations with just a linear error decay. In contrast to that, `HM-IRLS` exhibits a superlinear error decay quite early (i.e., for example as early as after 2 or 3 iterations), at least if the sample complexity is large enough, cf. Figure 1a.

Secondly, it can be observed that the convergence of the algorithm of [13] to a sparse vector often breaks down if p is smaller than 0.5 [13, Section 8]. In contrast to that, we observe that `HM-IRLS` does not suffer from this loss of global convergence for $p \ll 0.5$. Thus, a choice of very small parameters $p \approx 0.1$ or smaller is suggested as such a choice is accompanied by a very fast convergence, cf. Figure 3.

5.3 Recovery performance comparison with state-of-the-art algorithms

After comparing the performance of `HM-IRLS` with other IRLS variants, we now conduct experiments to compare `HM-IRLS` also to low-rank matrix recovery algorithms different from IRLS.

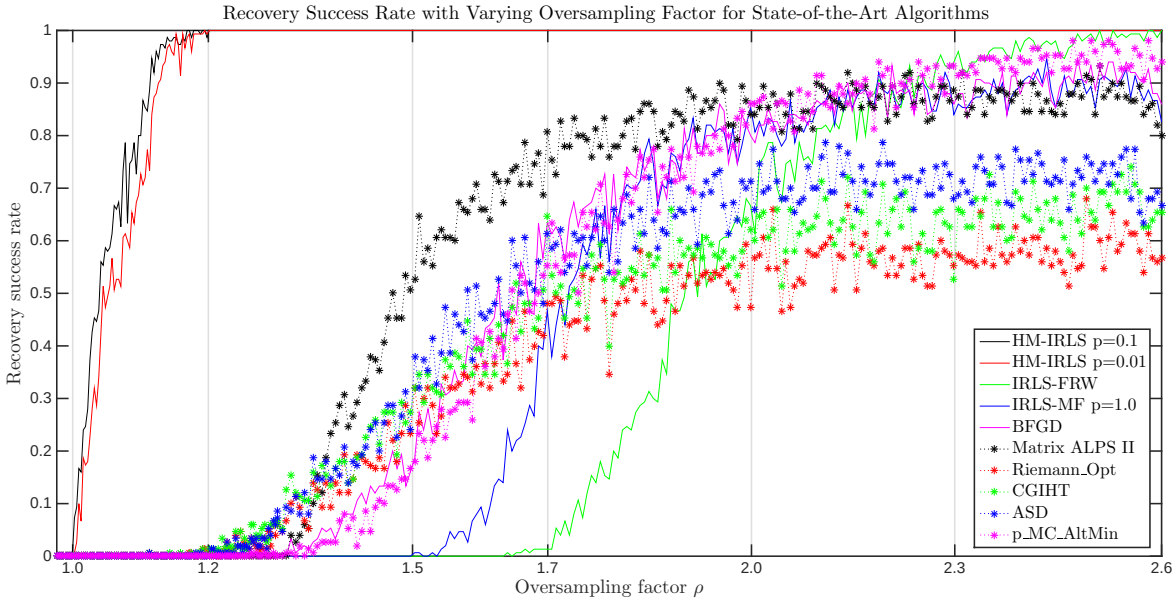

Figure 4: Recovery success rate with varying oversampling factor ρ for state-of-the-art-algorithms

To obtain a comprehensive picture, we consider a variety of state-of-the-art methods in the experiments, namely the already discussed IRLS algorithms `IRLS-FRW`, `IRLS-MF` [19, 39], but also the Riemannian optimization technique `Riemann_Opt` of [54], alternating minimization approaches `p_MC_AltMin`, `ASD` and `BFGD` of [28, 51, 41], and finally algorithms based on iterative hard thresholding such as `MatrixALPSII`, `CGIHT_Matrix` in [34, 4].

In the experiments, we examine the empirical recovery probabilities of the different algorithms systematically for varying sample complexities $m = \lfloor \rho \cdot d_f \rfloor$, parametrized by ρ, regulating the difficulty of the low-rank recovery problem. We recall that a large ρ describes an easy reconstruction problem, and a small parameter ρ, e.g., $\rho \approx 1$ describes very hard problems.

We choose $d_1 = d_2 = 100$ and the $r = 8$ as parameter of the experimental setting, conducting the experiments to recover rank-8 matrices $X_0 \in \mathbb{R}^{100\times 100}$. We remain in the matrix completion measurement setting described in subsection 5.1, but sample now 150 random instances of X_0 and Φ for different numbers of measurements varying between $m_{\min} = 1500$ to $m_{\max} = 4000$. This means that the oversampling factor ρ increases from $\rho_{\min} = 0.975$ to $\rho_{\max} = 2.60$. For each algorithm, a successful recovery of X_0 is defined as a relative Frobenius error $\|X^{\text{out}} - X_0\|_F / \|X_0\|_F$ of the matrix X^{out} returned by the algorithm of smaller than 10^{-3}. The algorithms are run until convergence or until a certain maximal number of iterations $n_{\max}$ (e.g. $n_{\max} = 10000$) is reached. The number $n_{\max}$ is chosen large enough to ensure that a recovery failure is not due to a lack of iterations.

In the experiments, we used implementations provided by the authors of the corresponding papers for the comparison algorithms, using default input parameters provided by the authors. The respective code sources can be found in the references.

5.3.1 Beyond the state-of-the-art performance of `HM-IRLS`

The results of the experiment can be seen in Figure 4. We observe that `HM-IRLS` exhibits a very high empirical recovery probability for $p = 0.1$ and $p = 0.01$ as soon as the sample complexity parameter ρ is larger than 1.0, which means that $m = \lfloor \rho r(d_1 + d_2 - r) \rfloor$ measurements suffice to recover $(d_1 \times d_2)$-dimensional rank-r matrices with ρ close to 1. This is very close to the information theoretical lower

bound of $d_f = r(d_1 + d_2 - r)$. Very interestingly, it can be observed that the empirical recovery percentage reaches 100% already for an oversampling factor of $\rho \approx 1.15$.

This is in stark contrast to the behavior of all the other algorithms that we considered in our experiments, which basically never recover any rank-r matrix if $\rho < 1.2$, and most of which need a sample complexity parameter of $\rho > 1.7$ to exceed a empirical recovery probability of a mere 50%. A success rate of close to 80% is reached not before raising ρ above 2.0 in our experimental setting, and also only for a subset of the comparison algorithms, in particular `Matrix ALPS II`, `BFGD`, `p_MC_AltMin` and the previous IRLS approaches `IRLS-FRW` and `IRLS-MF`. The empirical probability of 100%, as for `HM-IRLS`, is not systematically reached in our experimental setting for any competing method, even for quite large oversampling factors such as $\rho = 2.5$. While we do not rule out that a possible parameter tuning could improve the performance of any of the algorithms slightly, we conclude that for hard matrix completion problems, the experimental evidence for the vast differences in the recovery performance of `HM-IRLS` compared to other methods is striking.

Thus, our observation is that the proposed `HM-IRLS` algorithm *recovers low-rank matrices systematically with nearly the optimal number of measurements and needs fewer measurements than all the state-of-the-art algorithms we included in our experiments*, if the non-convexity parameter p is chosen such that $p \ll 1$.

We also note that the very sharp phase transition between failure and success that can be observed in Figure 4 for `HM-IRLS` indicates that the sample complexity parameter ρ is indeed the major variable determining the success of `HM-IRLS`. In contrast, the wider phase transitions for the other algorithms suggest that they might depend more on other factors, as the realizations of the random sampling model and the interplay of measurement operator Φ and ground truth matrix X_0.

Another conclusion that can be drawn from the empirical recovery probability of 100% is that, despite the severe non-convexity of the underlying Schatten-p quasinorm for, e.g., $p = 0.01$, `HM-IRLS` does not get stuck in local minimizers, but always converges to the global minimum if the oversampling factor is large enough. This suggests that for `HM-IRLS`, the initialization of $X^{(1)}$ as the Frobenius norm minimizer already ensures global convergence, while other non-convex methods might depend heavily on the initialization. In this sense, the results of our experiments can be interpreted such that competing non-convex low-rank matrix recovery algorithms seem to be not able to capture the desired basin of attraction of the global minimum in many cases if only few measurements are provided (i.e., if ρ is small), leading to the discovery of local minimzers only.

5.4 Computational complexity

While the harmonic mean weight matrix $\widetilde{W}^{(n)}$ (cf. (16)) is an inverse of a $(d_1 d_2 \times d_1 d_2)$-matrix and therefore in general a dense $(d_1 d_2 \times d_1 d_2)$-matrix, it is important to note that it never has to be computed explicitly in an implementation of `HM-IRLS`; neither is it necessary to compute its inverse $(\widetilde{W}^{(n)})^{-1} = \frac{1}{2}\left(U^{(n)}(\bar{\Sigma}^{(n)})^{2-p}U^{(n)*} \oplus V^{(n)}(\bar{\Sigma}^{(n)})^{2-p}V^{(n)*}\right)$ explicitly.

Indeed, as it can be seen in (10) and by the definition of the Kronecker sum (2), the harmonic mean weight matrix appears just as the linear operator $(\mathcal{W}^{(n)})^{-1}$ (cf. (13)) on the space of matrices $M_{d_1 \times d_2}$, whose action consists of a left- and right-sided matrix multiplication. Therefore, the application of $(\mathcal{W}^{(n)})^{-1}$ is $O(d_1 d_2 (d_1 + d_2))$ by the naive matrix multiplication algorithm, and can be easily parallelized.

The cost to compute $\Phi \circ \widetilde{\mathcal{W}}^{(n)-1} \circ \Phi^* \in M_{m \times m}$ depends on the linear measurement operator Φ. In the matrix completion setting (23), no additional arithmetic operations have to be performed, as Φ is a just a selection operator in this case.

Thus, the algorithm `HM-IRLS` consists of basically of two computational steps per iteration: The computation of the SVD of the $d_1 \times d_2$-matrix $X^{(n)}$ and the solution of the linearly constrained least squares problem in (14). The first is of time complexity $O(d_1 d_2 \min(d_1, d_2))$. The time complexity of the second depends on Φ, but is dominated by the inversion of a symmetric, $m \times m$ sparse linear system in the matrix completion setting, if m is the number of given entries. This has a time complexity of

at most $O(\max(d_1,d_2)^3 r^3)$ if ρ is just a constant oversampling factor.

For the matrix completion case, this allows us to recover low-rank matrices up to, e.g., $d_1 = d_2 = 1000$ on a single machine given very few entries.

Acceleration possibilities and extensions

To tackle higher dimensionalities in reasonable runtimes, a key strategy could be to address the computational bottleneck of `HM-IRLS`, the solution of the $m \times m$ linear system in (14), by using iterative methods. For IRLS algorithms designed for the related sparse recovery problem, the usage of conjugate gradient (CG) methods is discussed in [18]. By coupling the accuracy of the CG solutions to the outer IRLS iteration and using appropriate preconditioning, the authors obtain a competitive solver for the sparse recovery problem, also providing a convergence analysis. Similar ideas could be used for an acceleration of `HM-IRLS`.

It is interesting to see if further computational improvements can be achieved by combining the ideas of `HM-IRLS` with the usage of truncated and randomized SVDs [29], replacing the full SVDs of the $X^{(n)}$ that are needed to define the linear operator $(\mathcal{W}^{(n)})^{-1}$ in Algorithm 1.

6 Theoretical analysis

For the theoretical analysis of `HM-IRLS`, we introduce the following auxiliary functional $\mathcal{J}_p$, leading to a variational interpretation of the algorithm. In the whole section, we denote $d = \min(d_1,d_2)$ and $D = \max(d_1,d_2)$.

Definition 1. *Let $0 < p \leqslant 1$. Given a full-rank matrix $W \in M_{d_1 \times d_2}$, let*

$$\widetilde{W}(W) := 2\big[\mathbf{I}_{d_2} \otimes (WW^*)^{\frac{1}{2}}\big] \Big[(WW^*)^{\frac{1}{2}} \oplus (W^*W)^{\frac{1}{2}}\Big]^{-1} \big[(W^*W)^{\frac{1}{2}} \otimes \mathbf{I}_{d_1}\big] \in H_{d_1 d_2 \times d_1 d_2} \tag{24}$$

be the harmonic mean matrix $\widetilde{W}$ associated to W.

We define the auxiliary functional $\mathcal{J}_p : M_{d_1\times d_2} \times \mathbb{R}_{\geqslant 0} \times M_{d_1\times d_2} \to \mathbb{R}_{\geqslant 0}$ *as*

$$\mathcal{J}_p(X,\epsilon,W) := \begin{cases} \frac{p}{2}\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} + \frac{\epsilon^2 p}{2}\sum_{i=1}^{d}\sigma_i(W) + \frac{2-p}{2}\sum_{i=1}^{d}\sigma_i(W)^{\frac{p}{(p-2)}} & \text{if } \operatorname{rank}(W) = d, \\ +\infty & \text{if } \operatorname{rank}(W) < d. \end{cases} \tag{25}$$

We note that the matrix $\widetilde{W}$ of (24) is just the harmonic mean of the matrices $\widetilde{W}_1 := \mathbf{I}_{d_2} \otimes (WW^*)^{\frac{1}{2}}$ and $\widetilde{W}_2 = (W^*W)^{\frac{1}{2}} \otimes \mathbf{I}_{d_1}$, as introduced in section 2.4, if $(WW^*)^{\frac{1}{2}}$ *and* $(W^*W)^{\frac{1}{2}}$ are positive definite. Indeed, in this case, $(WW^*)^{\frac{1}{2}} \oplus (W^*W)^{\frac{1}{2}} = \widetilde{W}_1 + \widetilde{W}_2$ is invertible and as $(A^{-1}+B^{-1})^{-1} = A(A+B)^{-1}B$ for any positive definite matrices A and B of the same dimensions,

$$\widetilde{W}(W) = 2\widetilde{W}_1\big(\widetilde{W}_1 + \widetilde{W}_2\big)^{-1}\widetilde{W}_2 = 2(\widetilde{W}_1^{-1} + \widetilde{W}_2^{-1})^{-1}. \tag{26}$$

We use the more general definition (24) as it is well-defined for any full-rank $W \in M_{d_1\times d_2}$ and as it allows to handle the case of non-square matrices, i.e., the case $d_1 \neq d_2$, as in this case $(WW^*)^{\frac{1}{2}}$ or $(W^*W)^{\frac{1}{2}}$ has to be singular. Using the Moore-Penrose pseudo inverse $\widetilde{W}_1^+$ and $\widetilde{W}_2^+$ of the matrices $\widetilde{W}_1$ and $\widetilde{W}_2$, we can rewrite (24) as

$$\widetilde{W}(W) = 2\widetilde{W}_1\big(\widetilde{W}_1 + \widetilde{W}_2\big)^{-1}\widetilde{W}_2 = 2(\widetilde{W}_1^{+} + \widetilde{W}_2^{+})^{-1}.$$

With the auxiliary functional $\mathcal{J}_p$ at hand, we can interpret Algorithm 1 as an alternating minimization of the functional $\mathcal{J}_p(X,\epsilon,W)$ with respect to its arguments X, ϵ and W.

In the following, we derive the formula (16) for the weight matrix $\widetilde{W}^{(n+1)}$ as the evaluation $\widetilde{W}^{(n+1)} = \widetilde{W}\big(W^{(n+1)}\big)$ of $\widetilde{W}$ from Definition 1 at the minimizer

$$W^{(n+1)} = \operatorname*{arg\,min}_{W \in M_{d_1 \times d_2}} \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W), \tag{27}$$

with the minimizer being unique. Similarly, the formula (14) can be interpreted as

$$X^{(n+1)} = \operatorname*{arg\,min}_{\substack{X \in M_{d_1 \times d_2} \\ \Phi(X) = Y}} \|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W^{(n)}))} = \operatorname*{arg\,min}_{\substack{X \in M_{d_1 \times d_2} \\ \Phi(X) = Y}} \mathcal{J}_p(X, \epsilon^{(n)}, W^{(n)}) \tag{28}$$

These observations constitute the starting point of the convergence analysis of Algorithm 1, which is detailed subsequently after the verification of the optimization steps.

Remark 3. *Notice that* $\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} \neq \|X\|^2_{F((W*W)^{1/2})} \neq \|X^*\|^2_{F((WW*)^{1/2})}$. *The weighted norm on the left incorporates a symmetrized weighting acting in both the column* and *row space, while the ones on the right only reweight one-sided in the column* or *row space respectively.*

We claim above that the update rules for the iterates $X^{(n)}$ and $W^{(n)}$ in the steps (14) and (16) of Algorithm 1 are resulting from alternating minimizations of the functional $\mathcal{J}_p(X, \epsilon, W)$ with respect to the variables X, ϵ and W. In the upcoming subsection we will verify these assertions.

6.1 Optimization of $\mathcal{J}_p$ with respect to W

We fix $X \in M_{d_1 \times d_2}$ with singular value decomposition $X = \sum_{i=1}^{d} s_i u_i v_i^*$, where $u_i \in \mathbb{C}^{d_1}$ $v_i \in \mathbb{C}^{d_2}$ are the left and right singular vectors respectively and $s_i = s_i(X)$ denote its singular values for $i \in [d]$.

Our objective in the following is the justification of formula (16). To yield the building blocks of the weight matrix $\widetilde{W}^{(n+1)}$, we consider the minimization problem

$$\operatorname*{arg\,min}_{W \in M_{d_1 \times d_2}} \mathcal{J}_p(X, \epsilon, W) \tag{29}$$

for $\epsilon > 0$.

Lemma 4. *The unique minimizer of* (29) *is given by*

$$W_{\mathrm{opt}} = \sum_{i=1}^{d} (s_i(X)^2 + \epsilon^2)^{\frac{p-2}{2}} u_i v_i^*.$$

Furthermore, the value of $\mathcal{J}_p$ *at the minimizer* W_{opt} *is*

$$\mathcal{J}_p(X, \epsilon, W_{\mathrm{opt}}) = \sum_{i=1}^{d} (s_i(X)^2 + \epsilon^2)^{\frac{p}{2}} =: g_\epsilon^p(X) \tag{30}$$

for $p > 0$.

Remark 4. *We note that the value of* $\mathcal{J}_p(X, \epsilon, W_{\mathrm{opt}})$ *can be interpreted as a smooth* ϵ*-perturbation of a p-th power of a Schatten-p quasi-norm of the matrix* X. *In fact, for* $\epsilon = 0$ *we have*

$$\mathcal{J}_p(X, 0, W_{\mathrm{opt}}) = \|X\|_{S_p}^p = g_0^p(X).$$

Proof. First, we define the function

$$f^p_{X,\epsilon}(W) = \mathcal{J}_p(X, \epsilon, W) = \begin{cases} \frac{p}{2}\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} + \frac{\epsilon^2 p}{2} \sum_{i=1}^{d} \sigma_i(W) + \frac{2-p}{2} \sum_{i=1}^{d} \sigma_i(W)^{\frac{p}{(p-2)}} & \text{if } \operatorname{rank}(W) = d, \\ +\infty & \text{if } \operatorname{rank}(W) < d, \end{cases}$$

for $X \in M_{d_1 \times d_2}$, $\epsilon > 0$ fixed and with $W \in M_{d_1 \times d_2}$ as its only argument. We note that the set of minimizers of $f^p_{X,\epsilon}(W)$ does not contain an instance W with rank smaller than d as the value of $f^p_{X,\epsilon}(W)$ is infinite at such points and, therefore, it is sufficient to search for minimizers on the set $\Omega = \{Z \in M_{d_1 \times d_2} | \operatorname{rank}(Z) = d\}$ of matrices with rank d. We observe that the set Ω is an open set and that we have that

(a) $f^p_{X,\epsilon}(W)$ is lower semicontinuous, which means that any sequence $(W^k)_{k\in\mathbb{N}}$ with $W^k \overset{k\to\infty}{\longrightarrow} W$ fulfills $\liminf\limits_{k\to\infty} f^p_{X,\epsilon}(W^k) \geqslant f^p_{X,\epsilon}(W)$,

(b) $f^p_{X,\epsilon}(W) \geqslant \alpha$ for all $W \in M_{d_1 \times d_2}$ for some constant α,

(c) $f^p_{X,\epsilon}(W)$ is coercive, i.e., for any sequence $(W^k)_{k\in\mathbb{N}}$ with $\|W^k\|_F \overset{k\to\infty}{\longrightarrow} \infty$, we have $f^p_{X,\epsilon}(W^k) \overset{k\to\infty}{\longrightarrow} \infty$.

Property (a) is true as $f^p_{X,\epsilon}(W))|_\Omega$ is a concatenation of an indicator function of an open set, which is lower-semicontinuous and a sum of continuous functions on Ω. Property (b) is obviously true for the choice $\alpha = 0$.

To justify point (c), we note that $f^p_{X,\epsilon}(W) > \frac{\epsilon^2 p}{2} \sum\limits_{i=1}^{d} \sigma_i(W) = \frac{\epsilon^2 p}{2}\|W\|_{S_1} \geqslant \frac{\epsilon^2 p}{2}\|W\|_F$ and therefore, coercivity is clear from its definition. As a consequence from (a) and (c), it is also true that the level sets $L_C = \left\{W \in M_{d_1 \times d_2} | f^p_{X,\epsilon}(W) \leqslant C\right\}$ are closed and bounded and therefore, compact.

Via the direct method of calculus of variations, we conclude from the properties (a) - (c) that $f^p_{X,\epsilon}(W)$ has at least one global minimizer belonging to the set of critical points of $f^p_{X,\epsilon}(W)$ [12, Theorem 1].

To characterize the set of critical points of $f^p_{X,\epsilon}(W)$, its derivative with respect to W is calculated explicitly and equated with zero in the Appendix A. The solution of the resulting equation reveals that

$$W_{\text{opt}} = \sum_{i=1}^{d} (s_i^2 + \epsilon^2)^{\frac{p-2}{2}} u_i v_i^* =: \sum_{i=1}^{d} \sigma_i u_i v_i^*$$

is the only critical point and consequently the unique global minimizer of $f^p_{X,\epsilon}(W)$. We define the matrices $W^L_{\text{opt}} := \sum_{i=1}^{d} \sigma_i u_i u_i^*$ and $W^R_{\text{opt}} := \sum_{i=1}^{d} \sigma_i v_i v_i^*$, and note that

$$\widetilde{W}(W_{\text{opt}}) = 2(W_{\text{opt}}^{R^{-1}} \oplus W_{\text{opt}}^{L^{-1}})^{-1}$$

with Definition 1. To verify the second part of the theorem, we simply plug the optimal solution W_{opt} into the functional $\mathcal{J}_p$ and compute using (3) that

$$\begin{aligned}
\mathcal{J}_p(X, \epsilon, W_{\text{opt}}) &= \frac{p}{2}\|X_{\text{vec}}\|^2_{\ell_2(\widetilde{W}(W_{\text{opt}}))} + \frac{\epsilon^2 p}{2} \sum_{i=1}^{d} \sigma_i(W_{\text{opt}}) + \frac{2-p}{2} \sum_{i=1}^{d} \sigma_i(W_{\text{opt}})^{\frac{p}{p-2}} \\
&= \frac{p}{2} \sum_{i=1}^{d} \left[s_i^2 (u_i^* \otimes v_i^*) 2 \left(\sum_{k=1}^{d_2} \sum_{j=1}^{d_1} \frac{u_k u_k^* \otimes v_j v_j^*}{\sigma_k^{-1} + \sigma_j^{-1}} \right) (u_i \otimes v_i) \right]_{ii} + \frac{\epsilon^2 p}{2} \sum_{i=1}^{d} \sigma_i + \frac{2-p}{2} \sum_{i=1}^{d} \sigma_i^{\frac{p}{p-2}} \\
&= \frac{p}{2} \sum_{i=1}^{d} \left[2 s_i^2 \left(\sum_{k=1}^{d_2} \sum_{j=1}^{d_1} \frac{u_i^* u_k u_k^* u_i \otimes v_i^* v_j v_j^* v_i}{\sigma_k^{-1} + \sigma_j^{-1}} \right) \right]_{ii} + \frac{\epsilon^2 p}{2} \sum_{i=1}^{d} \sigma_i + \frac{2-p}{2} \sum_{i=1}^{d} \sigma_i^{\frac{p}{p-2}} \\
&= \frac{p}{2} \sum_{i=1}^{d} (s_i^2 + \epsilon^2)\sigma_i + \frac{2-p}{2} \sum_{i=1}^{d} \sigma_i^{\frac{p}{p-2}} = \frac{p}{2} \sum_{i=1}^{d} (s_i^2 + \epsilon^2)(s_i^2 + \epsilon^2)^{\frac{p-2}{2}} + \frac{2-p}{2} \sum_{i=1}^{d} (s_i^2 + \epsilon^2)^{\frac{p}{2}} \\
&= \sum_{i=1}^{d} (s_i^2 + \epsilon^2)^{\frac{p}{2}}.
\end{aligned}$$

□

6.2 Optimization of $\mathcal{J}_p$ with respect to X

Now, we prove the fact that our definition rule (14) of $X^{(n+1)}$ in Algorithm 1 can be interpreted as a minimization of the auxiliary functional $\mathcal{J}_p$ with respect to the variable X. Additionally, we show that this minimization step can be formulated as the solution of a weighted least squares problem with weight matrix $\widetilde{W}^{(n)}$. This is summarized in the following lemma.

Lemma 5. *Let $0 < p \leqslant 1$. Given a full-rank matrix $W \in M_{d_1 \times d_2}$, let $\widetilde{W}(W) := 2([(WW^*)^{\frac{1}{2}}]^+ \oplus [(W^*W)^{\frac{1}{2}}]^+)^{-1} \in H_{d_1 d_2 \times d_1 d_2}$ be the matrix from Definition 1 and $\mathcal{W}^{-1} : M_{d_1 \times d_2} \to M_{d_1 \times d_2}$ the linear operator of its inverse*

$$\mathcal{W}^{-1}(Z) := \frac{1}{2}\left[[(WW^*)^{\frac{1}{2}}]^+ Z + Z[(W^*W)^{\frac{1}{2}}]^+\right].$$

Then the matrix

$$X_{\mathrm{opt}} = \left(\mathcal{W}^{-1} \circ \Phi^* \circ (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)^{-1}\right)(Y) \in M_{d_1 \times d_2}$$

is the unique minimizer of the optimization problems

$$\operatorname*{arg\,min}_{\Phi(X)=Y} \mathcal{J}_p(X, \epsilon, W) = \operatorname*{arg\,min}_{\Phi(X)=Y} \|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W})}. \tag{31}$$

Moreover, a matrix $X_{\mathrm{opt}} \in M_{d_1 \times d_2}$ is a minimizer of the minimization problem (31) if and only if it fulfills the property

$$\langle \widetilde{W}(W)(X_{\mathrm{opt}})_{\mathrm{vec}}, H_{\mathrm{vec}} \rangle_{\ell_2} = 0 \quad \text{for all} \quad H \in \mathcal{N}(\Phi) \quad \text{and} \quad \Phi(X_{\mathrm{opt}}) = Y. \tag{32}$$

Proof. The equality of the optimization problems (31) can easily be seen by the fact that only the first summand of $\mathcal{J}_p(X, \epsilon, W)$ depends on X.

Next, we note that $\widetilde{W}(W) = 2([(W^*W)^{\frac{1}{2}}]^+ \oplus [(WW^*)^{\frac{1}{2}}]^+)^{-1}$ is positive definite: Let $W = \sum_{i=1}^d \sigma_i u_i v_i^*$, where u_i, v_i for $i \in [d]$ are the left and right singular vector respectively and σ_i for $i \in [d]$ are the singular values of W. Since $W^*W = \sum_{i=1}^d \sigma_i^2 v_i v_i^* \succeq 0$, also the generalized inverse root fulfills $[(WW^*)^{\frac{1}{2}}]^+ \succeq 0$ and for $WW^* = \sum_{i=1}^d \sigma_i^2 u_i u_i^* \succeq 0$, it follows that $[(WW^*)^{\frac{1}{2}}]^+ \succeq 0$. We stress that at least one of the matrices $(WW^*)^{\frac{1}{2}}$ and $(W^*W)^{\frac{1}{2}}$ is positive definite and therefore, $\frac{1}{2}[(W^*W)^{\frac{1}{2}}]^+ \oplus [(WW^*)^{\frac{1}{2}}]^+ = \frac{1}{2}[(W^*W)^{\frac{1}{2}}]^+ \otimes \mathbf{I}_{d_1} + \mathbf{I}_{d_2} \otimes [(WW^*)^{\frac{1}{2}}]^+ \succ 0$. This implies that also $\widetilde{W}(W) \succ 0$, as this is the inverse of the former matrix.

With the fact that $\widetilde{W}(W) \succ 0$, it can be proven analogously to the standard equivalence of [21, Proposition A.23] that a matrix $\bar{X} \in M_{d_1 \times d_2}$ is a minimizer of $\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} = \langle \widetilde{W}(W) X_{\mathrm{vec}}, X_{\mathrm{vec}} \rangle_{\ell_2}$ under the constraint $\Phi(X) = Y$ if and only if

$$\mathfrak{Re}\, \langle \widetilde{W}(W) \bar{X}_{\mathrm{vec}}, \psi_{\mathrm{vec}} \rangle_{\ell_2} = 0 \ \text{ for all } \ \psi \in \mathcal{N}(\Phi) \quad \text{and} \quad \Phi(\bar{X}) = Y.$$

Consequently, the optimal complex solution X_{opt} needs to fulfill the condition

$$\langle \widetilde{W}(W)(X_{\mathrm{opt}})_{\mathrm{vec}}, \psi_{\mathrm{vec}} \rangle_{\ell_2} = 0 \ \text{ for all } \ \psi \in \mathcal{N}(\Phi) \quad \text{and} \quad \Phi(X_{\mathrm{opt}}) = Y \tag{33}$$

as the derivative of the expression $\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))}$ with respect to a complex variable X has to satisfy the Cauchy-Riemann equations. Furthermore, the latter property holds if and only if $X_{\mathrm{opt}} = \left(\mathcal{W}^{-1} \circ \Phi^* \circ (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)^{-1}\right)(Y) \in M_{d_1 \times d_2}$. Indeed, (32) is equivalent to the existence of $\lambda \in \mathbb{C}^m$ such that $\Phi^*(\lambda)_{\mathrm{vec}} = \widetilde{W}(W)(X_{\mathrm{opt}})_{\mathrm{vec}}$. From the definition of $\widetilde{W}(W)^{-1}$ we see that $\widetilde{W}(W)^{-1} = \frac{1}{2}([(WW^*)^{\frac{1}{2}}]^+ \oplus [(W^*W)^{\frac{1}{2}}]^+)$ and the latter equation implies

$$X_{\mathrm{opt}} = \left[\widetilde{W}(W)^{-1}(\Phi^*(\lambda)_{\mathrm{vec}})\right]_{\mathrm{mat}} = \frac{1}{2}\left([(WW^*)^{1/2}]^+ \Phi^*(\lambda) + \Phi^*(\lambda)[(W^*W)^{1/2}]^+\right) = (\mathcal{W}^{-1} \circ \Phi^*)(\lambda). \tag{34}$$

As $Y = \Phi(X_{\mathrm{opt}}) = (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)(\lambda)$ in this case, we can calculate $\lambda = (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)^{-1}(Y)$ and we obtain from (34) that $X_{\mathrm{opt}} = (\mathcal{W}^{-1} \circ \Phi^* \circ (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)^{-1})(Y)$. On the other hand, any $X_{\mathrm{opt}} = (\mathcal{W}^{-1} \circ \Phi^* \circ (\Phi \circ \mathcal{W}^{-1} \circ \Phi^*)^{-1})(Y)$ fulfills (32) by construction. We conclude the proof by noting that X_{opt} fulfills the condition (32).

□

6.3 Basic properties of the algorithm

In the following subsection, we will have a closer look at Algorithm 1 and point out some of its properties, in particular, the boundedness of the iterates $(X^{(n)})_{n\in\mathbb{N}}$ and the fact that two consecutive iterates are getting arbitrarily close as $n \to \infty$. These results will be useful to develop finally the proof of convergence and to determine the rate of convergence of Algorithm 1 under conditions determined along the way.

First, we collect some observations about the functional $\mathcal{J}_p$ that appear analogously in the existing IRLS literature in the following lemma.

Lemma 6. *Let $(X^{(n)}, \epsilon^{(n)})_{n\in\mathbb{N}}$ be the sequence of iterates and smoothing parameters of Algorithm 1. Let $X^{(n)} = \sum_{i=1}^{d} \sigma_i^{(n)} u_i^{(n)} v_i^{(n)*}$ be the SVD of the n-th iterate $X^{(n)}$. Let $(W^{(n)})_{n\in\mathbb{N}}$ be a corresponding sequence such that*

$$W^{(n)} = \sum_{i=1}^{d} (\sigma_i^{(n)2} + \epsilon^{(n)2})^{\frac{p-2}{2}} u_i^{(n)} v_i^{(n)*}$$

for $n \in \mathbb{N}$. Then the following properties hold:

(a) $\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) \geqslant \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n+1)})$ for all $n \geqslant 1$,

(b) $\|X^{(n)}\|_{S_p}^p \leqslant \mathcal{J}_p(X^{(1)}, \epsilon^{(0)}, W^{(0)}) =: \mathcal{J}_{p,0}$ for all $n \geqslant 1$,

(c) The iterates $X^{(n)}, X^{(n+1)}$ come arbitrarily close as $n \to \infty$, i.e., $\lim_{n\to\infty} \|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|_{\ell_2}^2 = 0$.

Proof. (a) With the minimization property that defines $X^{(n+1)}$ in (28) and the inequality $\epsilon^{(n+1)} \leqslant \epsilon^{(n)}$, it follows that

$$\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) \geqslant \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n)}, W^{(n)}) \geqslant \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n)})$$

as $\mathcal{J}_p$ is monotonously increasing in ϵ. Furthermore, we obtain by the minimization property that defines $W^{(n+1)}$ in (27) and Lemma 4 that

$$\mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n)}) \geqslant \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n+1)}),$$

which finishes the proof of (a).

(b) For all $n \in \mathbb{N}$, we see that

$$\|X^{(n)}\|_{S_p}^p \leqslant g_{\epsilon^{(n)}}^p(X^{(n)}) = \mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) \leqslant \mathcal{J}_p(X^{(1)}, \epsilon^{(0)}, W^{(0)}),$$

using Lemma 4 and the monotonicity property of (a).

(c) For each $n \in \mathbb{N}$, we have with (a) and Definition 1 that

$$\begin{aligned}
&\frac{2}{p}\left[\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) - \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n+1)})\right] \\
&\geqslant \frac{2}{p}\left[\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) - \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n)}, W^{(n)}))\right] \\
&= \|X_{\mathrm{vec}}^{(n)}\|_{\ell_2(\widetilde{W}(W^{(n)}))}^2 - \|X_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde{W}(W^{(n)}))}^2 \\
&= \langle (X^{(n)} + X^{(n+1)})_{\mathrm{vec}}, (X^{(n)} - X^{(n+1)})_{\mathrm{vec}} \rangle_{\ell_2(\widetilde{W}^{(n)})},
\end{aligned}$$

using the notation $\widetilde{W}^{(n)} := \widetilde{W}(W^{(n)})$ in the last equality. As $X^{(n+1)}$ is the minimizer of $\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}^{(n)})}$ under the linear constraint and $X^{(n)} - X^{(n+1)} \in \mathcal{N}(\Phi)$, we deduce from (32) that

$$\langle \widetilde{W}^{(n)} X^{(n+1)}_{\mathrm{vec}}, (X^{(n)} - X^{(n+1)})_{\mathrm{vec}} \rangle = 0.$$

As a consequence, we get

$$\begin{aligned}\langle (X^{(n)} + X^{(n+1)})_{\mathrm{vec}}, (X^{(n)} - X^{(n+1)})_{\mathrm{vec}} \rangle_{\ell_2(\widetilde{W}^{(n)})} &= \langle (X^{(n)} - X^{(n+1)})_{\mathrm{vec}}, (X^{(n)} - X^{(n+1)})_{\mathrm{vec}} \rangle_{\ell_2(\widetilde{W}^{(n)})} \\ &= \|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}^{(n)})}.\end{aligned}$$

Moreover, we need an estimate on $\sigma_{d_1 d_2}(\widetilde{W}^{(n)})$ to obtain a bound on the difference of iterates independent of the weighting matrix. Since $1 = \sigma_{d_1 d_2}(Z)\sigma_1(Z^{-1})$ for any invertible matrix $Z \in M_{d_1 d_2 \times d_1 d_2}$, it is sufficient to compute $\sigma_1\big((\widetilde{W}^{(n)})^{-1}\big)$ to gain information on $\sigma_{d_1 d_2}\big(\widetilde{W}^{(n)}\big)$.

As the spectrum of a Kronecker sum $A \oplus B$ consists of the pairwise sum of the spectra of A and B [3, Proposition 7.2.3], we calculate that

$$\begin{aligned}\sigma_1^p\big((\widetilde{W}^{(n)})^{-1}\big) &= \left[\frac{1}{2}\Big((\sigma_1^{(n)2} + \epsilon^{(n)2})^{\frac{2-p}{2}} + (\sigma_1^{(n)2} + \epsilon^{(n)2})^{\frac{2-p}{2}}\Big)\right]^p = \big(\sigma_1^{(n)2} + \epsilon^{(n)2}\big)^{\frac{p}{2}(2-p)} \\ &\leqslant \Big(g^p_{\epsilon^{(n)}}(X^{(n)})\Big)^{2-p} = \Big(\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)})\Big)^{2-p} \leqslant \Big(\mathcal{J}_p(X^{(1)}, \epsilon^{(0)}, W^{(0)})\Big)^{2-p},\end{aligned}$$

using Lemma 4 and the monotonicity of $\mathcal{J}_p$ from Lemma 6.

Hence we conclude that

$$\sigma_{\min}(\widetilde{W}^{(n)}) = \sigma_{d_1 d_2}(\widetilde{W}^{(n)}) \geqslant \big(\mathcal{J}_p(X^{(1)}, \epsilon^{(0)}, W^{(0)})\big)^{1-\frac{2}{p}} = \mathcal{J}_{p,0}^{1-\frac{2}{p}}$$

and combine this with the previous calculations to obtain

$$\begin{aligned}&\frac{2}{p}\Big[\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) - \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n+1)})\Big] = \frac{2}{p}\|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}^{(n)})} \\ &\geqslant \frac{2}{p}\sigma_{\min}(\widetilde{W}^{(n)})\|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2} \geqslant \frac{2}{p}\mathcal{J}_{p,0}^{1-\frac{2}{p}}\|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2} \\ &= C_p\|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2}\end{aligned}$$

with the constant $C_p := \frac{2}{p}\mathcal{J}_{p,0}^{1-\frac{2}{p}} > 0$. From the monotonicity as in Lemma 6(a) and the boundedness of the sequence $\big(\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)})\big)_{n\in\mathbb{N}}$, we know that

$$\lim_{n\to\infty}\Big[\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) - \mathcal{J}_p(X^{(n+1)}, \epsilon^{(n+1)}, W^{(n+1)})\Big] = 0,$$

hence also

$$\lim_{n\to\infty}\|(X^{(n)} - X^{(n+1)})_{\mathrm{vec}}\|^2_{\ell_2} = 0.$$

□

At this point we notice that, assuming $X^{(n)} \to \bar{X}$ and $\epsilon^{(n)} \to \overline{\epsilon}$ for $n \to \infty$ with the limit point $(\bar{X}, \overline{\epsilon}) \in M_{d_1\times d_2} \times \mathbb{R}_{\geqslant 0}$, it would follow that

$$\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) \to g^p_{\overline{\epsilon}}(\bar{X})$$

for $n \to \infty$ by equation (30).

Now, let $\epsilon > 0$, a measurement vector $Y \in \mathbb{C}^m$ and the linear operator Φ be given and consider the optimization problem

$$\min_{\substack{X \in M_{d_1 \times d_2} \\ \Phi(X) = Y}} g_\epsilon^p(X) \tag{35}$$

with $g_\epsilon^p(X) = \sum_{i=1}^d (\sigma_i(X)^2 + \epsilon^2)^{\frac{p}{2}}$ and $\sigma_i(X)$ being the i-th singular value of X, cf. (30). If $g_\epsilon^p(X)$ is non-convex, which is the case for $p < 1$, one might practically only be able to find critical points of the problem.

Lemma 7. *Let $X \in M_{d_1 \times d_2}$ be a matrix with the SVD such that $X = \sum_{i=1}^d \sigma_i u_i v_i^*$, let $\epsilon > 0$. If we define*

$$\widetilde{W}(X, \epsilon) = 2\Big[\Big(\sum_{i=1}^d (\sigma_i^2 + \epsilon^2)^{\frac{2-p}{2}} u_i u_i^* \Big) \oplus \Big(\sum_{i=1}^d (\sigma_i^2 + \epsilon^2)^{\frac{2-p}{2}} v_i v_i^* \Big) \Big]^{-1} \in H_{d_1 d_2 \times d_1 d_2},$$

then $\widetilde{W}(X^{(n)}, \epsilon^{(n)}) = \widetilde{W}^{(n)}$, with $\widetilde{W}^{(n)}$ defined as in Algorithm 1, cf. (11).

Furthermore, X is a critical point of the optimization problem (35) *if and only if*

$$\langle \widetilde{W}(X, \epsilon) X_{\mathrm{vec}}, H_{\mathrm{vec}} \rangle_{\ell_2} = 0 \quad \text{for all} \quad H \in \mathcal{N}(\Phi) \quad \text{and} \quad \Phi(X) = Y. \tag{36}$$

In the case that g_ϵ^p is convex, i.e., if $p = 1$, (36) *implies that X is the unique minimizer of* (35).

Proof. The first statement $\widetilde{W}(X^{(n)}, \epsilon^{(n)}) = \widetilde{W}^{(n)}$ is clear from the definition of $\widetilde{W}(X, \epsilon)$ and (11).

To show the necessity of (36), let $X \in M_{d_1 \times d_2}$ be a critical point of (35). Without loss of generality, let us assume that $d_1 \leqslant d_2$. In this case, a short calculation shows that $g_\epsilon^p(X) = \operatorname{tr}\big[(XX^* + \epsilon^2 \mathbf{I}_{d_1})^{p/2}\big]$. It follows from the matrix derivative rules of [37, (7),(15),(18),(20) of Chapter 8.2] that

$$\nabla g_\epsilon^p(X) = p(XX^* + \epsilon^2 \mathbf{I}_{d_1})^{\frac{p-2}{2}} X = p \sum_{i=1}^d (\sigma_i^2 + \epsilon^2)^{\frac{p-2}{2}} \sigma_i u_i v_i^*,$$

using the singular value decomposition $X = \sum_{i=1}^d \sigma_i u_i v_i^*$ in the last equality.

Using the Kronecker sum inversion formula (3), we see that

$$\Big[\widetilde{W}(X, \epsilon) X_{\mathrm{vec}} \Big]_{\mathrm{mat}} = \sum_{i=1}^{d_1} \sum_{j=1}^{d_2} \frac{2\sigma_k}{(\sigma_i^2 + \epsilon^2)^{\frac{2-p}{2}} + (\sigma_j^2 + \epsilon^2)^{\frac{2-p}{2}}} \sum_{k=1}^d u_i u_i^* u_k v_k^* v_j v_j^* = \sum_{i=1}^d (\sigma_i^2 + \epsilon^2)^{\frac{p-2}{2}} \sigma_i u_i v_i^*$$

and thus $\nabla g_\epsilon^p(X) = p\big[\widetilde{W}(X, \epsilon) X_{\mathrm{vec}}\big]_{\mathrm{mat}}$.

Let $H \in \mathcal{N}(\Phi)$ be arbitrary and consider the function $h_\epsilon^p(t) = g_\epsilon^p(X + tH) - g_\epsilon^p(X)$. We note that $(h_\epsilon^p)'(0) = \langle \nabla g_\epsilon^p(X), H \rangle_F$. If X is a critical point of (35), then 0 is also a critical point of h_ϵ^p, i.e., $(h_\epsilon^p)'(0) = 0$, and therefore

$$0 = (h_\epsilon^p)'(0) = \langle \nabla g_\epsilon^p(X), H \rangle_F = p \big\langle \widetilde{W}(X, \epsilon) X_{\mathrm{vec}}, H_{\mathrm{vec}} \big\rangle_{\ell_2},$$

which implies (36).

For the sufficiency of (36), let $X \in M_{d_1 \times d_2}$ such that $\Phi(X) = Y$ and

$$\langle \widetilde{W}(X, \epsilon) X_{\mathrm{vec}}, H_{\mathrm{vec}} \rangle = 0$$

for all $H \in \mathcal{N}(\Phi)$. Using the calculation above, we obtain

$$0 = \langle \widetilde{W}(X, \epsilon) X_{\mathrm{vec}}, H_{\mathrm{vec}} \rangle = \frac{1}{p} \langle \nabla g_\epsilon^p(X), H \rangle_F,$$

which means that $\nabla g_\epsilon^p(X)$ is perpendicular to the null space $\mathcal{N}(\Phi)$ of Φ and hence

$$\nabla g_\epsilon^p(X) \in \mathrm{Ran}(\Phi^*) \text{ and } \Phi(X) = Y.$$

This are just the first order optimality conditions of (35) and therefore it follows that X is a critical point of g_ϵ^p under the linear constraint.

For $p = 1$, g_ϵ^p is strictly convex since $\epsilon > 0$. Thus, the problem (35) has a unique minimizer. Assuming that $X \in M_{d_1 \times d_2}$ fulfills (36), it follows that this minimizer is just X, as for any $X' \in M_{d_1 \times d_2}$ such that $\Phi(X') = Y$, $X - X' \in \mathcal{N}(\Phi)$ and therefore by convexity of g_ϵ^p,

$$g_\epsilon^p(X') \geqslant g_\epsilon^p(X) + \langle \nabla g_\epsilon^p(X), X' - X \rangle_F = g_\epsilon^p(X) + \langle \widetilde{W}(X, \epsilon) X_{\mathrm{vec}}, (X' - X)_{\mathrm{vec}} \rangle_{\ell_2} = g_\epsilon^p(X).$$

□

6.4 Convergence results under a null space property

We already presented the basic properties of the algorithm, as boundedness of the iterates of `HM-IRLS`, which are valid in any case. In this section, we provide convergence guarantees for the algorithm to stationary points of a smoothed Schatten-p functional g_ϵ^p as in (30) without any additional assumptions. Beyond that, assuming the measurement operator fulfills a Schatten-p null space property, we show the a-posteriori exact recovery statement that `HM-IRLS` converges to the low-rank matrix X_0 if $\lim_{n \to \infty} \epsilon_n = 0$, and a local convergence guarantee stating that the algorithms recovers the low-rank matrix X_0 if we obtain an iterate $X^{(\bar{n})}$ that is close enough to X_0.

Null space properties have already been used to guarantee the success of nuclear norm minimization

$$\min \|X\|_{S_1} \text{ subject to } \Phi(X) = Y,$$

or Schatten-1 minimization in our terminology, for solving the low-rank matrix recovery problem (5) [44]. To analyze the non-convex `HM-IRLS` algorithm, we introduce the following, related strong Schatten-p null space property [19, 40, 21].

Definition 2 (Strong Schatten-p null space property)**.** *Let $0 < p \leqslant 1$. We say that a linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ matrix fulfills the strong Schatten-p null space property (Schatten-p NSP) of order r with constant $0 < \gamma_r \leqslant 1$ if*

$$\Big(\sum_{i=1}^{r} \sigma_i^2(Z) \Big)^{p/2} < \frac{\gamma_r}{r^{1-\frac{p}{2}}} \Big(\sum_{i=r+1}^{d} \sigma_i^p(Z) \Big) \tag{37}$$

for all $Z \in \mathcal{N}(\Phi) \backslash \{0\}$.

We note that the definitions of Schatten-p null space properties are quite analogous to the ℓ_p-null space property in classical compressed sensing [21, Theorem 4.9], applied to the vector of singular values. In particular, (37) implies that

$$\sum_{i=1}^{r} \sigma_i^p(Z) < \sum_{i=r+1}^{d} \sigma_i^p(Z) \qquad \text{for all } Z \in \mathcal{N}(\Phi) \backslash \{0\}, \tag{38}$$

since $\|X\|_{S_p} \leqslant r^{1/p - 1/2} \|X\|_{S_2}$ for X that is rank-r. This, in turn, is a necessary and sufficient condition for solutions to (5) to be rank-r matrices.

Theorem 3 ([20])**.** *Let $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ be a linear map, let $0 < p \leqslant 1$ and $r \in \mathbb{N}$. Then every matrix $X_0 \in M_{d_1 \times d_2}$ with $\mathrm{rank}(X_0) \leqslant r$ and $\Phi(X_0) = Y \in \mathbb{C}^m$ can be exactly recovered by Schatten-p minimization* (7) *if and only if Φ fulfills* (38)*.*

Remark 5. *The sufficiency of the Schatten-p NSP* (38) *in Theorem 3 already been pointed out in [40]. The necessity as stated in the theorem, however, is due to a recent generalization of Mirsky's singular value inequalities to concave functions [2, 20].*

It can be seen that the (weak) Schatten-p NSP of (38) is a *stronger* property for larger p in the sense that if $0 < p' \leqslant p \leqslant 1$, the Schatten-$p$ property implies the Schatten-p' property. Very related to that, it can be seen that for any $0 < p \leqslant 1$, the strong Schatten-p null space property is implied by a sufficiently small *rank restricted isometry constant* δ_r, which is a classical tool in the analysis of low-rank matrix recovery algorithms [43, 7].

Definition 3. *The restricted isometry constant $\delta_r > 0$ of order r of the linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ is defined as the smallest number such that*

$$(1 - \delta_r)\|X\|_F^2 \leqslant \|\Phi(X)\|_{\ell_2}^2 \leqslant (1 + \delta_r)\|X\|_F^2$$

for all matrices $X \in M_{d_1 \times d_2}$ of rank at most r.

Indeed, it follows from the proof of [11, Theorem 4.1] that a restricted isometry constant of order $2r$ such that $\delta_{2r} < \frac{2}{\sqrt{2}+3} \approx 0.4531$ implies the strong Schatten-p NSP of order r with a constant $\gamma_r < 1$ for any $0 < p \leqslant 1$. More precisely, it can be seen that $\delta_{2r} < \frac{2}{\sqrt{2}+3}$ implies that the strong Schatten-p NSP (37) of order r holds with the constant $\gamma_r = \frac{(\sqrt{2}+1)^p}{2^p} \frac{\delta_{2r}^p}{(1-\delta_{2r})^p}$.

As the restricted isometry constant of order $2r$ is small for many random measurement models, e.g. if Φ is Gaussian as in (17) or more generally, sub-Gaussian, in the optimal measurement regime with overwhelming probability [7], the following lemma can therefore be deduced immediately.

Lemma 8. *For any $0 < p \leqslant 1$, $0 < \gamma < 1$ and any (sub-)Gaussian random operator $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ (e.g. as defined in* (17)*), there exist constants $C_1 > 1$, $C_2 > 0$ such that if $m \geqslant C_1(r(d_1 + d_2))$, the strong Schatten-p null space property* (37) *of order r with constant $\gamma_r < \gamma$ is fulfilled with probability at least $1 - e^{-C_2 m}$.*

For the convergence analysis of Algorithm 1, we use the following *reverse triangle inequalities* implied by the strong Schatten-p NSP:

Lemma 9. *Let $0 < p \leqslant 1$. Assume that the linear map $\Phi : M_{d_1 \times d_2} \to \mathbb{C}^m$ fulfills the strong Schatten-p-NSP* (37) *of order r with constant $\gamma_r \in (0, 1)$. Recall the definition of the best rank-r Schatten-p approximation error*

$$\beta_r(Z)_{S_p} := \inf\left\{\|Z - \widetilde{Z}\|_{S_p}^p,\ \widetilde{Z} \in M_{d_1 \times d_2} \text{ has rank } r\right\} = \sum_{i=r+1}^{d} \sigma_i(Z)^p \tag{39}$$

of a matrix $Z \in M_{d_1 \times d_2}$.
Let $Z, Z' \in M_{d_1 \times d_2}$ such that $\Phi(Z - Z') = 0$. Then

$$\|Z' - Z\|_F^p \leqslant \frac{2^p \gamma_r^{1-p/2}}{r^{1-p/2}} \frac{1}{1 - \gamma_r} \left(\|Z'\|_{S_p}^p - \|Z\|_{S_p}^p + 2\beta_r(Z)_{S_p}\right). \tag{40}$$

Proof. The proof of Lemma 9 is an adaptation of the proof of the corresponding result for ℓ_p-minimization [22, Theorem 13] using the generalization of Mirksy's singular value inequality to concave functions [2, 20]. Furthermore, the proof of the similar statement [33, Threorem 12] can be adapted to show (40). □

Remark 6. *However, if $m < d_1 d_2$, null space property-type assumptions as* (37) *or* (38) *do not hold for the important case of matrix completion-type measurements [8], where $\Phi(X)$ is given as m sample entries*

$$\Phi(X)_\ell = X_{i_\ell, j_\ell}, \qquad \ell = 1, \dots, m, \tag{41}$$

and $(i_\ell, j_\ell) \in [d_1] \times [d_2]$ for all $\ell \in [m]$, of the matrix $X \in M_{d_1 \times d_2}$.

Finally, we prove the convergence result Theorem 4 for `HM-IRLS`, making use of the Schatten-p-null space property discussed above. Having established Theorem 4, Theorem 1 follows as a corollary.

Theorem 4. *Let $\Phi : M_{d_1\times d_2} \to \mathbb{C}^m$ be a linear operator $Y \in \mathrm{Ran}(\Phi)$ a vector in its range. Let $(X^{(n)})_{n\geqslant 1}$ and $(\epsilon^{(n)})_{n\geqslant 1}$ be the sequences produced by Algorithm 1 for input parameters Φ, Y, r and $0 < p \leqslant 1$, let $\epsilon = \lim_{n\to\infty} \epsilon^{(n)}$.*

(i) If $\epsilon = 0$ and if Φ fulfills the strong Schatten-p NSP (37) of order r with constant $0 < \gamma_r < 1$, then the sequence $(X^{(n)})_{n\geqslant 1}$ converges to a matrix $\bar{X} \in M_{d_1\times d_2}$ of rank at most r that is the unique minimizer of the Schatten-p minimization problem (7). Moreover, there exists an absolute constant $\hat{C} > 0$ such that for any Z with $\Phi(Z) = Y$ and any $\widetilde{r} \leqslant r$, it holds that

$$\|Z - \bar{X}\|_F^p \leqslant \frac{\hat{C}}{r^{1-p/2}} \beta_{\widetilde{r}}(Z)_{S_p}, \tag{42}$$

where $\hat{C} = \frac{2^{p+1}\gamma_r^{1-p/2}}{1-\gamma_r}$ and $\beta_{\widetilde{r}}(Z)_{S_p}$ is the best rank-$\widetilde{r}$ Schatten-p approximation error, cf. (39).

(ii) If $\epsilon > 0$, then each accumulation point $\bar{X}$ of $(X^{(n)})_{n\geqslant 1}$ is a stationary point of the ϵ-perturbed Schatten-p functional g_ϵ^p of (18) under the linear constraint $\Phi(X) = Y$. If additionally $p = 1$, then $\bar{X}$ is the unique global minimizer of g_ϵ^p.

(iii) If Φ fulfills the strong Schatten-p NSP of order $2r$ with constant $\gamma_{2r} < 1$, assume that there exists a matrix $X_0 \in M_{d_1\times d_2}$ with $\Phi(X_0) = Y$ of rank $\widetilde{r} \leqslant r \leqslant \frac{\min(d_1,d_2)}{2}$, a constant $0 < \rho < 1$ and an iteration $\bar{n} \in \mathbb{N}$ such that

$$\|X^{(\bar{n})} - X_0\|_{S_\infty} \leqslant \rho\sigma_{\widetilde{r}}(X_0)$$

and $\epsilon^{\bar{n}} = \sigma_{r+1}(X^{\bar{n}})$.
If the condition number $\kappa = \frac{\sigma_1(X_0)}{\sigma_{\widetilde{r}}(X_0)}$ of X_0 and ρ are sufficiently small (see (45) and (46)), then

$$X^{(n)} \to X_0 \quad \text{for } n \to \infty.$$

Proof. (i) We first assume that there exists an iteration $\bar{n} \in \mathbb{N}$ such that $\epsilon^{(\bar{n})} = 0$. Then we define $\bar{X} := X^{(\bar{n})}$, which fulfills by construction that $\Phi(\bar{X}) = Y$ and $\sigma_{r+1}(\bar{X}) = 0$, i.e., $\mathrm{rank}(\bar{X}) \leqslant r$.

Otherwise, we have that $\epsilon^{(n)} > 0$ for all $n \in \mathbb{N}$. In this case, there exists a subsequence $(n_\ell)_{l\in\mathbb{N}}$ of $(n)_{n\geqslant n_0}$ such that $\epsilon^{(n_\ell)} < \epsilon^{(n_\ell-1)}$ for all $\ell \in \mathbb{N}$. By Lemma 6(b), $(X^{(n)})_n$ is bounded and we can extract a further subsequence, which we denote again by $(X^{(n_\ell)})_l$, that converges to some $\bar{X} := \lim_{l\to\infty} X^{(n_\ell)}$. Since $\lim_{l\to\infty} \epsilon^{(n_\ell)} = 0$, it follows as well that $\lim_{l\to\infty} \sigma_{r+1}(X^{(n_\ell)}) \leqslant \lim_{l\to\infty} \epsilon^{(n_\ell-1)} = 0$. Moreover, it follows by Weyl's stability estimate for the $(r+1)$-th singular value [19, Theorem 7.1] that $\sigma_{r+1}(\bar{X}) = 0$. Therefore, also for this case, $\bar{X}$ fulfills $\Phi(\bar{X}) = Y$ and $\mathrm{rank}(\bar{X}) \leqslant r$.

Now, we show that the whole sequence converges to $\bar{X}$. According to (30), it holds that

$$\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) = g^p_{\epsilon^{(n)}}(X^{(n)})$$

for any $n \in \mathbb{N}$. Since $(X^{(n_\ell)}) \xrightarrow{l\to\infty} \bar{X}$ and $\epsilon^{(n_\ell)} \xrightarrow{l\to\infty} 0$, we see therefore that

$$\mathcal{J}_p(X^{(n_\ell)}, \epsilon^{(n_\ell)}, W^{(n_\ell)}) \xrightarrow{l\to\infty} \|\bar{X}\|_{S_p}^p.$$

If we use the non-increasing monotonicity from Lemma 6(a), this also shows already that the same holds true for the whole sequence $(X^{(n)})_{n\geqslant 1}$, i.e., $\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) \xrightarrow{n\to\infty} \|\bar{X}\|_{S_p}^p$. Applying the triangle inequality for the p-power of the Schatten-p quasinorms $\|\cdot\|_{S_p}^p$, we obtain

$$\mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}) - d(\epsilon^{(n)})^p \leqslant \|X^{(n)}\|_{S_p}^p \leqslant \mathcal{J}_p(X^{(n)}, \epsilon^{(n)}, W^{(n)}).$$

Since $\lim_{n\to\infty} d(\epsilon^{(n)})^p = 0$, we conclude that also

$$\|X^{(n)}\|_{S_p}^p \xrightarrow{n\to\infty} \|\bar{X}\|_{S_p}^p.$$

Now it remains to show that $X^{(n)} \to \bar{X}$. Lemma 9 implies that

$$\|\bar{X} - X^{(n)}\|_F^p \leqslant \frac{2^p \gamma_r^{1-p/2}}{r^{1-p/2}} \frac{1}{1-\gamma_r} \left(\|\bar{X}\|_{S_p}^p - \|X^{(n)}\|_{S_p}^p + 2\beta_{r+1}(\bar{X})_{S_p} \right)$$

and therefore, with $\beta_{r+1}(\bar{X})_{S_p} = 0$, we obtain

$$\|\bar{X} - X^{(n)}\|_F^p \leqslant \frac{2^p \gamma_r^{1-p/2}}{r^{1-p/2}} \frac{1}{1-\gamma_r} \left(\|\bar{X}\|_{S_p}^p - \|X^{(n)}\|_{S_p}^p \right).$$

Hence, taking the limit $n \to \infty$ shows that $\lim\limits_{n\to\infty} \|\bar{X} - X^{(n)}\|_F = 0$, and consequently $X^{(n)} \to \bar{X}$ for $n \to \infty$.

As $\bar{X}$ is of rank at most r and also fulfills $\Phi(X) = Y$, the strong Schatten-p null space property implies via Lemma 9 that $\bar{X}$ is indeed the unique solution to (7) and that it coincides with X_0.

To obtain the error bound (42), we note that any matrix Z with $\Phi(Z) = Y$ fulfills

$$\|Z - \bar{X}\|_F^p \leqslant \frac{2^p \gamma_r^{1-p/2}}{r^{1-p/2}} \frac{1}{1-\gamma_r} \left(\|\bar{X}\|_{S_p}^p - \|Z\|_{S_p}^p + 2\beta_r(Z)_{S_p} \right) \leqslant \frac{2^{p+1} \gamma_r^{1-p/2}}{r^{1-p/2}} \frac{1}{1-\gamma_r} \beta_r(Z)_{S_p}, \qquad (43)$$

using Lemma 9 in the first inequality and the fact that $\bar{X}$ is the unique minimizer of (6) in the second inequality. The proof of (42) is then concluded by noting that $\beta_r(Z)_{S_p} \leqslant \beta_{\widetilde{r}}(Z)_{S_p}$ if $\widetilde{r} \leqslant r$.

(ii) We shall first show that $X^{(n)} \to X^\epsilon$, $n \to \infty$, with $X^\epsilon \in M_{d_1 \times d_2}$ being a stationary point of g_ϵ^p under the linear constraint. We already observed that $(X^{(n)})_{n\geqslant 1}$ is a bounded sequence and hence this sequence has accumulation points. Let $(X^{(n_\ell)})_{\ell\geqslant 1}$ be any convergent subsequence of $(X^{(n)})_{n\geqslant 1}$ and let $\bar{X}$ be its limit.

Since $\widetilde{W}(X, \epsilon)$ as defined in Lemma 7 depends continuously on X and ϵ, it follows that

$$\lim_{\ell\to\infty} \widetilde{W}(X^{(n_\ell)}, \epsilon^{n_\ell}) = \widetilde{W}(\bar{X}, \epsilon) =: \overline{W}.$$

On the other hand, by invoking Lemma 6(c), we also obtain that $X^{(n_\ell+1)} \to \bar{X}$ for $\ell \to \infty$ and therefore, also $\lim\limits_{\ell\to\infty} \widetilde{W}^{(n_\ell+1)} = \overline{W}$. We observe that from the interpretation of $X^{(n_\ell+1)}$ as a minimizer of $\mathcal{J}_p$ in (31) and the proof of Lemma 5, it follows that

$$\widetilde{W}^{(n_\ell)} X^{(n_\ell+1)} \in \mathrm{Ran}(\Phi^*) \text{ and } \Phi(X^{(n_\ell+1)}) = Y.$$

This implies that there exists $\lambda \in \mathbb{R}^m$ such that $\widetilde{W}^{(n_\ell+1)} X^{(n_\ell)} = \Phi^*(\lambda)$. We observe that for all $H \in \mathcal{N}(\Phi)$ and all $\ell \in \mathbb{N}$,

$$\langle \widetilde{W}^{(n_\ell)} X_{\mathrm{vec}}^{(n_\ell+1)}, H_{\mathrm{vec}} \rangle = \langle \Phi^*(\lambda), H_{\mathrm{vec}} \rangle = \langle \lambda, \Phi(\eta) \rangle = 0.$$

Consequently, $\langle \overline{W} \bar{X}_{\mathrm{vec}}, \eta_{\mathrm{vec}} \rangle = \lim\limits_{\ell\to\infty} \langle \widetilde{W}^{(n_\ell)} X_{\mathrm{vec}}^{(n_\ell+1)}, H_{\mathrm{vec}} \rangle = 0$ for all $H \in \mathcal{N}(\Phi)$. As shown in Lemma 7, this implies that $\bar{X}$ is a stationary point of g_ϵ^p and even coincides with the unique minimizer X^ϵ in the convex case, i.e., if $p = 1$.

(iii) This statement is a consequence of Theorem 5, which is proven below. □

Proof of Theorem 1(i-ii). For Theorem 1(i), it suffices to combine Theorem 4(i) with Lemma 8. Theorem 1(ii) follows trivially from Theorem 4(ii). □

6.5 Locally superlinear convergence

In this section, we show Theorem 5, which states that `HM-IRLS` exhibits a locally superlinear convergence rate if the operator Φ fulfills an appropriate Schatten-p null space property. Having proven this, Theorem 2 from section 4 follows then as a corollary.

Theorem 5 (Locally Superlinear Convergence Rate)**.** *Assume that the linear map $\Phi : M_{d_1\times d_2} \to \mathbb{C}^m$ fulfills the strong Schatten-p NSP of order $2r$ with constant $\gamma_{2r} < 1$ and that there exists a matrix $X_0 \in M_{d_1\times d_2}$ with $\operatorname{rank}(X_0) = r \leqslant \frac{\min(d_1,d_2)}{2}$ such that $\Phi(X_0) = Y$, let Φ, Y, r and $0 < p \leqslant 1$ be the input parameters of Algorithm 1. Moreover, let $\kappa = \frac{\sigma_1(X_0)}{\sigma_r(X_0)}$ be the condition number of X_0 and $\eta^{(n)} := X^{(n)} - X_0$ be the residuals of the n-th output matrix of Algorithm 1 for $n \in \mathbb{N}$.*
Assume that there exists an iteration $\bar{n} \in \mathbb{N}$ and a constant $0 < \rho < 1$ such that

$$\|\eta^{(\bar{n})}\|_{S_\infty} \leqslant \rho\sigma_r(X_0) \tag{44}$$

and $\epsilon^{(\bar{n})} = \sigma_{r+1}(X^{(\bar{n})})$.

If additionally the condition number κ and ρ are small enough such that

$$\mu\|\eta^{(\bar{n})}\|_{S_\infty}^{p(1-p)} < 1 \tag{45}$$

with the constant

$$\mu := 2^{5p}(1+\gamma_{2r})^p\Big(\frac{\gamma_{2r}(3+\gamma_{2r})(1+\gamma_{2r})}{(1-\gamma_{2r})}\Big)^{2-p}\Big(\frac{d-r}{r}\Big)^{2-\frac{p}{2}} r^p \frac{\sigma_r(X_0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p \tag{46}$$

then

$$\|\eta^{(n+1)}\|_{S_\infty} \leqslant \mu^{1/p}\left(\|\eta^{(n)}\|_{S_\infty}\right)^{2-p}$$

and

$$\|\eta^{(n+1)}\|_{S_p} \leqslant \mu^{1/p}\left(\|\eta^{(n)}\|_{S_p}\right)^{2-p}$$

for all $n \geqslant \bar{n}$.

In the proof of Theorem 5 we use the following bound on perturbations of the singular value decomposition, which is originally due to Wedin [55]. It bounds the alignment of the subspaces spanned by the singular vectors of two matrices by their norm distance, given a gap between the first singular values of the one matrix and the last singular values of the other matrix that is sufficiently pronounced.

Lemma 10 (Wedin's bound [47])**.** *Let Z and $\bar{Z}$ be two matrices of the same size and their singular value decompositions*

$$Z = \begin{pmatrix} U_1 & U_2\end{pmatrix}\begin{pmatrix}\Sigma_1 & 0\\ 0 & \Sigma_2\end{pmatrix}\begin{pmatrix}V_1^*\\ V_2^*\end{pmatrix} \quad \text{and} \quad \bar{Z} = \begin{pmatrix}\bar{U}_1 & \bar{U}_2\end{pmatrix}\begin{pmatrix}\bar{\Sigma}_1 & 0\\ 0 & \bar{\Sigma}_2\end{pmatrix}\begin{pmatrix}\bar{V}_1^*\\ \bar{V}_2^*\end{pmatrix},$$

where the submatrices have the sizes of corresponding dimensions. Suppose that δ, α satisfying $0 < \delta \leqslant \alpha$ are such that $\alpha \leqslant \sigma_{\min}(\Sigma_1)$ and $\sigma_{\max}(\bar{\Sigma}_2) < \alpha - \delta$. Then

$$\|\bar{U}_2^* U_1\|_{S_\infty} \leqslant \sqrt{2}\frac{\|Z-\bar{Z}\|_{S_\infty}}{\delta} \text{ and } \|\bar{V}_2^* V_1\|_{S_\infty} \leqslant \sqrt{2}\frac{\|Z-\bar{Z}\|_{S_\infty}}{\delta}. \tag{47}$$

As a first step towards the proof of Theorem 2, we show the following lemma.

Lemma 11. *Let $(X^{(n)})_n$ be the output sequence of Algorithm 1 for parameters Φ, Y, r and $0 < p \leqslant 1$, and $X_0 \in M_{d_1\times d_2}$ be a matrix such that $\Phi(X_0) = Y$.*

(i) Let $\eta_{2r}^{(n+1)}$ be the best rank-$2r$ approximation of $\eta^{(n+1)} = X^{(n+1)} - X_0$. Then

$$\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^{2p} \leqslant 2^{2-p}\Bigg(\sum_{i=r+1}^{d}\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}}\Bigg)^{2-p}\|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde{W}^{(n)})}^{2p},$$

where $\widetilde{W}^{(n)}$ denotes the harmonic mean weight matrix from (11).

(ii) Assume that the linear map $\Phi : \mathbb{R}^{d_1\times d_2} \to \mathbb{R}^m$ fulfills the strong Schatten-p NSP of order $2r$ with constant $\gamma_{2r} < 1$. Then

$$\|\eta^{(n+1)}\|_{S_2}^{2p} \leqslant 2^p\frac{\gamma_{2r}^{2-p}}{r^{2-p}}\Bigg(\sum_{i=r+1}^{d}\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}}\Bigg)^{2-p}\|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde{W}^{(n)})}^{2p} \tag{48}$$

(iii) Under the same assumption as for (ii), it holds that

$$\|\eta^{(n+1)}\|_{S_p}^{2p} \leqslant (1+\gamma_{2r})^2 2^{2-p}\Bigg(\sum_{i=r+1}^{d}\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}}\Bigg)^{2-p}\|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde{W}^{(n)})}^{2p}.$$

Proof of Lemma 11. (i) Let the $X^{(n)} = \widetilde{U}^{(n)}\Sigma^{(n)}\widetilde{V}^{(n)*}$ be the (full) singular value decomposition of $X^{(n)}$, i.e., $\widetilde{U}^{(n)} \in U_{d_1}$ and $\widetilde{V}^{(n)} \in U_{d_2}$ are unitary matrices and $\Sigma^{(n)} = \mathrm{diag}(\sigma_1(X^{(n)}), \dots, \sigma_r(X^{(n)})) \in M_{d_1\times d_2}$. We define $U_T^{(n)} \in U_{d_1\times r}$ as the matrix of the first r columns of $\widetilde{U}^{(n)}$ and $U_{T_c}^{(n)} \in U_{d_1\times(d_1-r)}$ as the matrix of its last $d_1 - r$ columns, so that $\widetilde{U}^{(n)} = \begin{pmatrix} U_T^{(n)} & U_{T_c}^{(n)}\end{pmatrix}$, and similarly $V_T^{(n)}$ and $V_{T_c}^{(n)}$.

As $\mathbf{I}_{d_1} = U_T^{(n)}U_T^{(n)*} + U_{T_c}^{(n)}U_{T_c}^{(n)*}$ and $\mathbf{I}_{d_2} = V_T^{(n)}V_T^{(n)*} + V_{T_c}^{(n)}V_{T_c}^{(n)*}$, we note that

$$U_{T_c}^{(n)}U_{T_c}^{(n)*}\eta^{(n+1)}V_{T_c}^{(n)}V_{T_c}^{(n)*} = \eta^{(n+1)} - U_T^{(n)}U_T^{(n)*}\eta^{(n+1)} + U_{T_c}^{(n)}U_{T_c}^{(n)*}\eta^{(n+1)}V_T^{(n)}V_T^{(n)*},$$

while $U_T^{(n)}U_T^{(n)*}\eta^{(n+1)} + U_{T_c}^{(n)}U_{T_c}^{(n)*}\eta^{(n+1)}V_T^{(n)}V_T^{(n)*}$ has a rank of at most $2r$. This implies that

$$\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p} \leqslant \|U_{T_c}^{(n)}U_{T_c}^{(n)*}\eta^{(n+1)}V_{T_c}^{(n)}V_{T_c}^{(n)*}\|_{S_p} = \|U_{T_c}^{(n)*}\eta^{(n+1)}V_{T_c}^{(n)}\|_{S_p}. \tag{49}$$

Using the definitions of $\widetilde{U}^{(n)}$ and $\widetilde{V}^{(n)}$, we write the harmonic mean weight matrices of the n-th iteration (11) as

$$\widetilde{W}^{(n)} = 2(\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})\big(\bar{\Sigma}_{d_1}^{(n)2-p} \oplus \bar{\Sigma}_{d_2}^{(n)2-p}\big)^{-1}(\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})^*, \tag{50}$$

where $\bar{\Sigma}_{d_1}^{(n)} \in M_{d_1\times d_1}$ and $\bar{\Sigma}_{d_2}^{(n)} \in M_{d_2\times d_2}$ are the diagonal matrices with the smoothed singular values of $X^{(n)}$ from (12), but filled up with zeros if necessary. Using the abbreviation

$$\Omega := (\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})^*\widetilde{W}^{(n)\frac{1}{2}}\eta_{\mathrm{vec}}^{(n+1)} \in \mathbb{C}^{d_1d_2}, \tag{51}$$

we rewrite

$$\begin{aligned}\eta_{\mathrm{vec}}^{(n+1)} &= \widetilde{W}^{(n)-\frac{1}{2}}\widetilde{W}^{(n)\frac{1}{2}}\eta_{\mathrm{vec}}^{(n+1)} = 2^{-1/2}(\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})\big(\bar{\Sigma}_{d_1}^{(n)2-p} \oplus \bar{\Sigma}_{d_2}^{(n)2-p}\big)^{1/2}\Omega \\ &= 2^{-1/2}(\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})\Big[(\mathbf{I}_{d_2} \otimes \bar{\Sigma}_{d_1}^{(n)\frac{2-p}{2}})D_L + (\bar{\Sigma}_{d_2}^{(n)\frac{2-p}{2}} \otimes \mathbf{I}_{d_1})D_R\Big]\Omega\end{aligned} \tag{52}$$

with the diagonal matrices $D_L, D_R \in M_{d_1d_2\times d_1d_2}$ such that

$$(D_L)_{i+(j-1)d_1,i+(j-1)d_1} = \Big(1 + \Big(\frac{\sigma_j^2(X^{(n)}) + \epsilon^{(n)2}}{\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}}\Big)^{\frac{2-p}{2}}\Big)^{-1/2}$$

and

$$(D_R)_{i+(j-1)d_1,i+(j-1)d_1} = \Big(\Big(\frac{\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}}{\sigma_j^2(X^{(n)}) + \epsilon^{(n)2}}\Big)^{\frac{2-p}{2}} + 1\Big)^{-1/2}$$

for $i \in [d_1]$ and $j \in [d_2]$. This can be seen from the definitions of the Kronecker product $\otimes$ and the Kronecker sum $\oplus$ (cf. section 2.2), as $\Big(\big(\bar\Sigma_{d_1}^{(n)2-p} \oplus \bar\Sigma_{d_2}^{(n)2-p}\big)^{1/2}\Big)_{i+(j-1)d_1,i+(j-1)d_1} = (s_i + s_j)^{1/2} = s_i(s_i+s_j)^{-1/2} + s_j(s_i+s_j)^{-1/2} = s_i^{1/2}(1+\frac{s_j}{s_i})^{-1/2} + s_j^{1/2}(\frac{s_i}{s_j}+1)^{-1/2}$, if s_ℓ denotes the ℓ-th diagonal entry of $\bar\Sigma_{d_2}^{(n)2-p}$ and $\bar\Sigma_{d_1}^{(n)2-p}$ for $\ell \in [\max(d_1,d_2)]$.

If we write $\bar\Sigma_{d_1,T_c}^{(n)\frac{2-p}{2}} \in M_{(d_1-r)\times(d_1-r)}$ for the diagonal matrix containing the $d_1 - r$ last diagonal elements of $\bar\Sigma_{d_1}^{(n)2-p}$ and $\bar\Sigma_{d_2,T_c}^{(n)\frac{2-p}{2}} \in M_{(d_1-r)\times(d_1-r)}$ for the diagonal matrix containing the $d_2 - r$ last diagonal elements of $\bar\Sigma_{d_2}^{(n)2-p}$, it follows from (52) that

$$\begin{aligned}
\big\|U_{T_c}^{(n)*}\eta^{(n+1)}V_{T_c}^{(n)}\big\|_{S_p}^p &= 2^{-\frac{p}{2}}\Big\|U_{T_c}^{(n)*}\widetilde U^{(n)}\Big[\bar\Sigma_{d_1}^{(n)\frac{2-p}{2}}(D_L\Omega)_{\mathrm{mat}} + (D_R\Omega)_{\mathrm{mat}}\bar\Sigma_{d_2}^{(n)\frac{2-p}{2}}\Big]\widetilde V^{(n)*}V_{T_c}^{(n)}\Big\|_{S_p}^p \\
&= 2^{-\frac{p}{2}}\Big\|\bar\Sigma_{d_1,T_c}^{(n)\frac{2-p}{2}}\big[(D_L\Omega)_{\mathrm{mat}}\big]_{T_c,T_c} + \big[(D_R\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\bar\Sigma_{d_2,T_c}^{(n)\frac{2-p}{2}}\Big\|_{S_p}^p \qquad (53)\\
&\leqslant 2^{-\frac{p}{2}}\Big\|\bar\Sigma_{d_1,T_c}^{(n)\frac{2-p}{2}}\big[(D_L\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\Big\|_{S_p}^p + \Big\|\big[(D_R\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\bar\Sigma_{d_2,T_c}^{(n)\frac{2-p}{2}}\Big\|_{S_p}^p
\end{aligned}$$

with the notation that M_{T_c,T_c} denotes the submatrix of M which contains the intersection of the last $d_1 - r$ rows of M with its last $d_2 - r$ columns.

Now, Hölder's inequality for Schatten-p quasinorms (e.g. [23, Theorem 11.2]) can be used to see that

$$\Big\|\bar\Sigma_{d_1,T_c}^{(n)\frac{2-p}{2}}\big[(D_L\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\Big\|_{S_p}^p \leqslant \Big\|\bar\Sigma_{T_c}^{(n)\frac{2-p}{2}}\Big\|_{S_{\frac{2p}{2-p}}}^p\Big\|\big[(D_L\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\Big\|_{S_2}^p. \qquad (54)$$

Inserting the definition

$$\Big\|\bar\Sigma_{T_c}^{(n)\frac{2-p}{2}}\Big\|_{S_{\frac{2p}{2-p}}}^p = \Bigg(\sum_{i=r+1}^d\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{2-p}{4}\frac{2p}{2-p}}\Bigg)^{\frac{2-p}{2}} = \Bigg(\sum_{i=r+1}^d\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}}\Bigg)^{\frac{2-p}{2}}$$

allows us to rewrite the first factor, while the second factor can be bounded by

$$\begin{aligned}
\big\|\big[(D_L\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\big\|_{S_2}^p &\leqslant \big\|(D_L\Omega)_{\mathrm{mat}}\big\|_{S_2}^p \leqslant \|\Omega_{\mathrm{mat}}\|_{S_2}^p = \|(\widetilde V^{(n)}\otimes\widetilde U^{(n)})^*\widetilde W^{(n)\frac12}\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2}^p \\
&= \|\widetilde W^{(n)\frac12}\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2}^p = \|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde W^{(n)})}^p,
\end{aligned}$$

as the matrix $D_L \in M_{d_1d_2\times d_1d_2}$ from (52) fulfills $\|D_L\|_{S_\infty} \leqslant 1$ since its entries are bounded by 1; we also recall the definition (51) of Ω and that $\widetilde V^{(n)}$ and $\widetilde U^{(n)}$ are unitary.

The term $\Big\|\big[(D_R\Omega)_{\mathrm{mat}}\big]_{T_c,T_c}\bar\Sigma_{d_2,T_c}^{(n)\frac{2-p}{2}}\Big\|_{S_p}^p$ in (53) can be estimated analogously. Combining this with (49), we obtain

$$\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^{2p} \leqslant 2^{-p}\Big(2\Big(\sum_{i=r+1}^d\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}}\Big)^{\frac{2-p}{2}}\Big)^2\|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde W^{(n)})}^{2p},$$

concluding the proof of statement (i).

(ii) Using the strong Schatten-p null space property (37) of order $2r$ and that $\eta^{(n+1)} \in \mathcal N(\Phi)$, we estimate

$$\begin{aligned}
\|\eta^{(n+1)}\|_{S_2}^{2p} &= \big(\|\eta_{2r}^{(n+1)}\|_{S_2}^2 + \|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_2}^2\big)^p \leqslant \Big(\frac{\gamma_{2r}^{2/p} + \gamma_{2r}^{2/p-1}}{(2r)^{2/p-1}}\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^2\Big)^p \\
&\leqslant \frac{\gamma_{2r}^{2-p}(\gamma_{2r}+1)^p}{(2r)^{2-p}}\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^{2p} \leqslant 2^p\frac{\gamma_{2r}^{2-p}}{2^{2-p}r^{2-p}}\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^{2p},
\end{aligned} \qquad (55)$$

where we use in the second inequality a version of Stechkin's lemma [33, Lemma 3.1], which leads to the estimate

$$\|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_2}^2 \leqslant \Big(\frac{\|\eta_{2r}^{(n+1)}\|_{S_2}}{2r}\Big)^{2-p} \|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^p \leqslant \frac{\gamma_{2r}^{2/p-1}}{(2r)^{2/p-1}} \|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^2.$$

Combining (55) with statement (i), this results in

$$\|\eta^{(n+1)}\|_{S_2}^{2p} \leqslant 2^p \frac{\gamma_{2r}^{2-p}}{r^{2-p}} \Bigg(\sum_{i=r+1}^{d} \big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{\frac{p}{2}} \Bigg)^{2-p} \|\eta_{\mathrm{vec}}^{(n+1)}\|_{\ell_2(\widetilde{W}^{(n)})}^{2p},$$

which shows statement (ii).

(iii) For the third statement, we use the strong Schatten-p NSP (37) to see that

$$\|\eta^{(n+1)}\|_{S_p}^p = \|\eta_{2r}^{(n+1)}\|_{S_p}^p + \|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^p \leqslant (1+\gamma_{2r}) \|\eta^{(n+1)} - \eta_{2r}^{(n+1)}\|_{S_p}^p,$$

and combine this with statement (i). □

Lemma 12. *Let $(X^{(n)})_n$ be the output sequence of Algorithm 1 with parameters Φ, Y, r and $0 < p \leqslant 1$, and $\widetilde{W}^{(n)}$ be the harmonic mean weight matrix matrix (11) for $n \in \mathbb{N}$. Let $X_0 \in M_{d_1 \times d_2}$ be a rank-r matrix such that $\Phi(X_0) = Y$ with condition number $\kappa := \frac{\sigma_1(X_0)}{\sigma_r(X_0)}$.*

(i) If (44) is fulfilled for iteration n, then $\eta^{(n+1)} = X^{(n)} - X_0$ fulfills

$$\Big\|\eta_{\mathrm{vec}}^{(n+1)}\Big\|_{\ell_2(\widetilde{W}^{(n)})}^{2p} \leqslant \frac{4^p r^{p/2} (s_r^0)^{p(p-1)}}{(1-\rho)^{2p}} \kappa^p \frac{\|\eta^{(n)}\|_{S_\infty}^{2p-p^2}}{(\epsilon^{(n)})^{2p-p^2}} \|\eta^{(n+1)}\|_{S_2}^p.$$

(ii) Under the same assumption as for (i), it holds that

$$\Big\|\eta_{\mathrm{vec}}^{(n+1)}\Big\|_{\ell_2(\widetilde{W}^{(n)})}^{2p} \leqslant \frac{7^p r^{p/2} \max(r, d-r)^{p/2} (s_r^0)^{p(p-1)}}{(1-\rho)^{2p}} \kappa^p \frac{\|\eta^{(n)}\|_{S_\infty}^{2p-p^2}}{(\epsilon^{(n)})^{2p-p^2}} \|\eta^{(n+1)}\|_{S_\infty}^p$$

Proof of Lemma 12. (i) Recall that $X^{(n+1)} = \operatorname*{arg\,min}_{\Phi(X)=Y} \|X_{\mathrm{vec}}\|_{\ell_2(\widetilde{W}^{(n)})}^2$ is the minimizer of the weighted least squares problem with weight matrix $\widetilde{W}^{(n)}$. As $\eta^{(n+1)} = X^{(n+1)} - X_0$ is in the null space of the measurement map Φ, it follows from Lemma 5 that

$$0 = \langle \widetilde{W}^{(n)} X_{\mathrm{vec}}^{(n+1)}, \eta_{\mathrm{vec}}^{(n+1)} \rangle = \langle \widetilde{W}^{(n)} (\eta^{(n+1)} + X_0)_{\mathrm{vec}}, \eta_{\mathrm{vec}}^{(n+1)} \rangle,$$

which is equivalent to

$$\Big\|\eta_{\mathrm{vec}}^{(n+1)}\Big\|_{\ell_2(\widetilde{W}^{(n)})}^2 = \langle \widetilde{W}^{(n)} \eta_{\mathrm{vec}}^{(n+1)}, \eta_{\mathrm{vec}}^{(n+1)} \rangle = -\langle \widetilde{W}^{(n)} (X_0)_{\mathrm{vec}}, \eta_{\mathrm{vec}}^{(n+1)} \rangle. \tag{56}$$

Using Hölder's inequality, we can therefore estimate

$$\begin{aligned} \Big\|\eta_{\mathrm{vec}}^{(n+1)}\Big\|_{\ell_2(\widetilde{W}^{(n)})}^2 &= -\langle \widetilde{W}^{(n)} (X_0)_{\mathrm{vec}}, \eta_{\mathrm{vec}}^{(n+1)} \rangle_{\ell_2} = -\langle \big[\widetilde{W}^{(n)} (X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}, \eta^{(n+1)} \rangle_F \\ &\leqslant \big\|\big[\widetilde{W}^{(n)} (X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\big\|_{S_2} \|\eta^{(n+1)}\|_{S_2}. \end{aligned} \tag{57}$$

To bound the first factor, we first rewrite the action of $\widetilde{W}^{(n)}$ on X_0 in the matrix space as

$$\begin{aligned} \Big[\widetilde{W}^{(n)} (X_0)_{\mathrm{vec}}\Big]_{\mathrm{mat}} &= 2[(\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)}) \big(\bar{\Sigma}_{d_1}^{(n)2-p} \oplus \bar{\Sigma}_{d_2}^{(n)2-p}\big)^{-1} (\widetilde{V}^{(n)} \otimes \widetilde{U}^{(n)})^* (X_0)_{\mathrm{vec}}]_{\mathrm{mat}} = \\ &= \widetilde{U}^{(n)} \big(H^{(n)} \circ (\widetilde{U}^{(n)*} X_0 \widetilde{V}^{(n)})\big) \widetilde{V}^{(n)*}, \end{aligned}$$

using (50) and Lemma 11 about the action of inverses of Kronecker sums, with the notation that $H^{(n)} \in M_{d_1 \times d_2}$ such that

$$H_{ij}^{(n)} = 2\Big[\mathbb{1}_{\{i \leqslant d\}}(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2})^{\frac{2-p}{2}} + \mathbb{1}_{\{j \leqslant d\}}(\sigma_j^2(X^{(n)}) + \epsilon^{(n)2})^{\frac{2-p}{2}}\Big]^{-1}$$

for $i \in [d_1]$, $j \in [d_2]$, where $\mathbb{1}_{\{i \leqslant d\}} = 1$ if $i \leqslant d$ and $\mathbb{1}_{\{i \leqslant d\}} = 0$ otherwise. This enables us to estimate

$$\begin{aligned}
\Big\|\big[\widetilde{W}^{(n)}(X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\Big\|_{S_2}^2 &= \Big\|\widetilde{U}^{(n)}\big(H^{(n)} \circ (\widetilde{U}^{(n)*} X_0 \widetilde{V}^{(n)})\big)\widetilde{V}^{(n)*}\Big\|_{S_2}^2 = \Big\|H^{(n)} \circ (\widetilde{U}^{(n)*} X_0 \widetilde{V}^{(n)})\Big\|_{S_2}^2 \\
&= \left\|H^{(n)} \circ \begin{pmatrix} U_T^{(n)*} X_0 V_T^{(n)} & U_T^{(n)*} X_0 V_{T_c}^{(n)} \\ U_{T_c}^{(n)*} X_0 V_T^{(n)} & U_{T_c}^{(n)*} X_0 V_{T_c}^{(n)} \end{pmatrix}\right\|_{S_2}^2 \\
&= \Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*} X_0 V_T^{(n)})\Big\|_{S_2}^2 + \Big\|H_{T,T_c}^{(n)} \circ (U_T^{(n)*} X_0 V_{T_c}^{(n)})\Big\|_{S_2}^2 \\
&\quad + \Big\|H_{T_c,T}^{(n)} \circ (U_{T_c}^{(n)*} X_0 V_T^{(n)})\Big\|_{S_2}^2 + \Big\|H_{T_c,T_c}^{(n)} \circ (U_{T_c}^{(n)*} X_0 V_{T_c}^{(n)})\Big\|_{S_2}^2,
\end{aligned} \tag{58}$$

using the notation from the proof of Lemma 11. To bound the first summand, we calculate

$$\begin{aligned}
&\Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*} X_0 V_T^{(n)})\Big\|_{S_2} \leqslant \Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*} X^{(n)} V_T^{(n)})\Big\|_{S_2} + \Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*}(-\eta^{(n)}) V_T^{(n)})\Big\|_{S_2} \\
&\leqslant \Big\|H_{T,T}^{(n)} \circ \Sigma_T^{(n)}\Big\|_{S_2} + \Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*} \eta^{(n)} V_T^{(n)})\Big\|_{S_2} \\
&\leqslant \left(\sum_{i=1}^{r} \frac{\sigma_i^2(X^{(n)})}{\big(\sigma_i^2(X^{(n)}) + \epsilon^{(n)2}\big)^{2-p}}\right)^{1/2} + \max_{i,j=1}^{r} |H_{i,j}^{(n)}| \|U_T^{(n)*} \eta^{(n)} V_T^{(n)}\|_{S_2} \\
&\leqslant \sqrt{r}\sigma_r^{p-1}(X^{(n)}) + (\sigma_r^2(X^{(n)}) + \epsilon^{(n)2})^{\frac{p-2}{2}} \|U_T^{(n)*} \eta^{(n)} V_T^{(n)}\|_{S_2} \\
&\leqslant \sqrt{r}\sigma_r^{p-1}(X^{(n)}) + \sigma_r^{p-2}(X^{(n)})\sqrt{r}\|\eta^{(n)}\|_{S_\infty} = \sqrt{r}\sigma_r^{p-2}(X^{(n)})\big[\sigma_r(X^{(n)}) + \|\eta^{(n)}\|_{S_\infty}\big],
\end{aligned} \tag{59}$$

denoting $\Sigma_T^{(n)} = \mathrm{diag}(\sigma_i(X^{(n)}))_{i=1}^r$ and that the matrices $U_T^{(n)}$ and $V_T^{(n)}$ contain the first r left resp. right singular vectors of $X^{(n)}$ in the second inequality, together with the estimates $\|X\|_{S_1} \leqslant \sqrt{r}\|X\|_{S_2} \leqslant r\|X\|_{S_\infty}$ for $(r \times r)$-matrices X.

We recall the notations $s_r^0 = \sigma_r(X_0)$ and $s_1^0 = \sigma_1(X_0)$ and note that

$$\sigma_r(X^{(n)}) \geqslant s_r^0(1-\rho),$$

as the assumption (44) implies that

$$s_r^0 = \sigma_r(X_0) = \sigma_r(X^{(n)} - \eta^{(n)}) \leqslant \sigma_r(X^{(n)}) + \sigma_1(\eta^{(n)}) \leqslant \sigma_r(X^{(n)}) + \rho s_r^0,$$

using [3, Proposition 9.6.8] in the first inequality.

Therefore we can bound the term of (59) such that

$$\Big\|H_{T,T}^{(n)} \circ (U_T^{(n)*} X_0 V_T^{(n)})\Big\|_{S_2} \leqslant \sqrt{r}(s_r^0(1-\rho))^{p-2}\big[s_r^0(1-\rho) + \rho s_r^0\big] = \sqrt{r}(s_r^0)^{p-1}(1-\rho)^{p-2}. \tag{60}$$

For the second summand in the estimate of $\Big\|\big[\widetilde{W}^{(n)}(X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\Big\|_{S_2}^2$, similar arguments and again as-

sumption (44) are used to compute

$$\begin{aligned}
&\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}X_0V_{T_c}^{(n)})\right\|_{S_2}\leqslant\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}X^{(n)}V_{T_c}^{(n)})\right\|_{S_2}+\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}\eta^{(n)}V_{T_c}^{(n)})\right\|_{S_2}\\
&=\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}U_T^{(n)}\Sigma_T^{(n)}\overbrace{V_T^{(n)*}V_{T_c}^{(n)}}^{=0}+\overbrace{U_T^{(n)*}U_{T_c}^{(n)}}^{=0}\Sigma_{T_c}^{(n)}V_{T_c}^{(n)*}V_{T_c}^{(n)})\right\|_{S_2}+\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}\eta^{(n)}V_{T_c}^{(n)})\right\|_{S_2}\\
&\leqslant\max_{i\in[r],j\in\{r+1,\dots,d_2\}}|H^{(n)}_{i,j}|\|U_T^{(n)*}\eta^{(n)}V_{T_c}^{(n)}\|_{S_2}\leqslant 2\big[(\sigma_r(X^{(n)})^2+\epsilon^{(n)2})^{\frac{2-p}{2}}\big]^{-1}\|U_T^{(n)*}\eta^{(n)}V_{T_c}^{(n)}\|_F\\
&\leqslant 2\sigma_r(X^{(n)})^{p-2}\|U_T^{(n)*}\eta^{(n)}V_{T_c}^{(n)}\|_{S_2}\leqslant 2\sqrt{r}(s_r^0(1-\rho))^{p-2}\|\eta^{(n)}\|_{S_\infty}\leqslant 2\rho\sqrt{r}(s_r^0)^{p-1}(1-\rho)^{p-2}.
\end{aligned}\tag{61}$$

From exactly the same arguments it follows that also

$$\left\|H^{(n)}_{T_c,T}\circ(U_{T_c}^{(n)*}X_0V_T^{(n)})\right\|_{S_2}\leqslant 2\rho\sqrt{r}(s_r^0)^{p-1}(1-\rho)^{p-2}.\tag{62}$$

It remains to bound the last summand $\left\|H^{(n)}_{T_c,T_c}\circ(U_{T_c}^{(n)*}X_0V_{T_c}^{(n)})\right\|_{S_2}^2$. We see that

$$\begin{aligned}
\left\|H^{(n)}_{T_c,T_c}\circ(U_{T_c}^{(n)*}X_0V_{T_c}^{(n)})\right\|_{S_2}&\leqslant\max_{\substack{i\in\{r+1,\dots,d_1\}\\ j\in\{r+1,\dots,d_2\}}}|H^{(n)}_{i,j}|\|U_{T_c}^{(n)*}X_0V_{T_c}^{(n)}\|_{S_2}\\
&\leqslant(\epsilon^{(n)})^{p-2}\|U_{T_c}^{(n)*}X_0V_{T_c}^{(n)}\|_{S_2}\leqslant(\epsilon^{(n)})^{p-2}\|U_{T_c}^{(n)*}U_T^0\|_{S_\infty}\|S^0\|_{S_2}\|V_T^{0*}V_{T_c}^{(n)}\|_{S_\infty}\\
&\leqslant(\epsilon^{(n)})^{p-2}\frac{\sqrt{2}\|\eta^{(n)}\|_{S_\infty}}{(1-\rho)s_r^0}\sqrt{r}s_1^0\frac{\sqrt{2}\|\eta^{(n)}\|_{S_\infty}}{(1-\rho)s_r^0}=2\sqrt{r}\|\eta^{(n)}\|_{S_\infty}^2(\epsilon^{(n)})^{p-2}(1-\rho)^{-2}(s_r^0)^{-1}\frac{s_1^0}{s_r^0}
\end{aligned}\tag{63}$$

where Hölder's inequality for Schatten norms was used in the third inequality. In the fourth inequality, Wedin's singular value perturbation bound of Lemma 10 is used with the choice $Z=X_0$, $\bar{Z}=X^{(n)}$, $\alpha=s_r^0$ and $\delta=(1-\rho)s_r^0$, and finally $\epsilon^{(n)}\leqslant\rho s_r^0$ in the last inequality, which is implied by the rule (15) for $\epsilon^{(n)}$ together with assumption (44).

Summarizing the estimates (60), (61), (62) and (63), we conclude that

$$\begin{aligned}
&\left\|\big[\widetilde{W}^{(n)}(X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\right\|_{S_2}^2\leqslant r(s_r^0)^{2p-2}(1-\rho)^{2p-4}\left[1+8\rho^2+4\frac{\|\eta^{(n)}\|_{S_\infty}^4}{(1-\rho)^{2p}}(\epsilon^{(n)})^{2p-4}(s_r^0)^{-2p}\left(\frac{s_1^0}{s_r^0}\right)^2\right]\\
&=\frac{r(s_r^0)^{2p-2}}{(1-\rho)^4}\left[(1+8\rho^2)(1-\rho)^{2p}+4\frac{\|\eta^{(n)}\|_{S_\infty}^{4-2p}}{(\epsilon^{(n)})^{4-2p}}\frac{\|\eta^{(n)}\|_{S_\infty}^{2p}}{(s_r^0)^{2p}}\left(\frac{s_1^0}{s_r^0}\right)^2\right]\\
&\leqslant\frac{r(s_r^0)^{2p-2}}{(1-\rho)^4}\left[9+4\frac{\|\eta^{(n)}\|_{S_\infty}^{4-2p}}{(\epsilon^{(n)})^{4-2p}}\rho^{2p}\kappa^2\right]\leqslant\frac{13r(s_r^0)^{2p-2}}{(1-\rho)^4}\left[\frac{\|\eta^{(n)}\|_{S_\infty}^{4-2p}}{(\epsilon^{(n)})^{4-2p}}\kappa^2\right],
\end{aligned}$$

as $0<\rho<1$, $\epsilon^{(n)}\leqslant\sigma_{r+1}(X^{(n)})=\|X^{(n)}_{T_c}\|_{S_\infty}\leqslant\|\eta^{(n)}\|_{S_\infty}$ and using the assumption (44) in the second inequality. This concludes the proof of Lemma 12(i) together with inequality (57) as $13^{p/2}\leqslant 16^{p/2}=4^p$.

(ii) For the second statement of Lemma 12, we proceed similarly as before, but note that by Hölder's inequality, also

$$\left\|\eta_{\mathrm{vec}}^{(n+1)}\right\|^2_{\ell_2(\widetilde{W}^{(n)})}\leqslant\big\|\big[\widetilde{W}^{(n)}(X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\big\|_{S_1}\|\eta^{(n+1)}\|_{S_\infty},$$

cf. (57). Furthermore

$$\begin{aligned}
\big\|\big[\widetilde{W}^{(n)}(X_0)_{\mathrm{vec}}\big]_{\mathrm{mat}}\big\|_{S_1}&\leqslant\left\|H^{(n)}_{T,T}\circ(U_T^{(n)*}X_0V_T^{(n)})\right\|_{S_1}+\left\|H^{(n)}_{T,T_c}\circ(U_T^{(n)*}X_0V_{T_c}^{(n)})\right\|_{S_1}\\
&\quad+\left\|H^{(n)}_{T_c,T}\circ(U_{T_c}^{(n)*}X_0V_T^{(n)})\right\|_{S_1}+\left\|H^{(n)}_{T_c,T_c}\circ(U_{T_c}^{(n)*}X_0V_{T_c}^{(n)})\right\|_{S_1}.
\end{aligned}$$

The four Schatten-1 norms can then be estimated by $\max(r,(d-r))^{1/2}$ times the corresponding Schatten-2 norms. Using then again inequalities (60) – (63), we conclude the proof of (ii). □

Proof of Theorem 5. First we note that

$$\left(\sum_{i=r+1}^{d}\left(\sigma_i^2(X^{(n)})+\epsilon^{(n)2}\right)^{\frac{p}{2}}\right)^{2-p} \leqslant 2^{p-\frac{p^2}{2}}(d-r)^{2-p}\sigma_{r+1}(X^{(n)})^{p(2-p)} \tag{64}$$

as $\epsilon^{(n)} \leqslant \sigma_{r+1}(X^{(n+1)})$ due to the choice of $\epsilon^{(n)}$ in (15). We proceed by induction over $n \geqslant \bar{n}$. Lemmas 11(ii) and 12(ii) imply together with (64) that for $n=\bar{n}$,

$$\begin{aligned}\|\eta^{(n+1)}\|_{S_\infty}^p &\leqslant \frac{\|\eta^{(n+1)}\|_{S_2}^{2p}}{\|\eta^{(n+1)}\|_{S_\infty}^p} \leqslant 2^p\gamma_{2r}^{2-p}2^{p-\frac{p^2}{2}}\Big(\frac{d-r}{r}\Big)^{2-p/2}\frac{7^p r^p (s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p\|\eta^{(n)}\|_{S_\infty}^{2p-p^2}\\ &\leqslant 2^{5p}\gamma_{2r}^{2-p}\Big(\frac{d-r}{r}\Big)^{2-p/2}\frac{r^p (s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p\|\eta^{(n)}\|_{S_\infty}^{p(2-p)}\end{aligned} \tag{65}$$

as $\sigma_{r+1}(X^{(n)})=\epsilon^{(n)}$ by assumption for $n=\bar{n}$.

Similarly, by Lemmas 11(iii) and 12(ii) and (64), the error in the Schatten-p quasinorm fulfills

$$\|\eta^{(n+1)}\|_{S_p}^{2p} \leqslant (1+\gamma_{2r})^2 2^{2+2p}\big(d-r\big)^{2-p}\frac{r^{p/2}(s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p\|\eta^{(n)}\|_{S_\infty}^{p(2-p)}\|\eta^{(n+1)}\|_{S_2}^p \tag{66}$$

for $n=\bar{n}$. Using the strong Schatten-p null space property of order $2r$ for the operator Φ, we see from the arguments of (55) that

$$\|\eta^{(n)}\|_{S_\infty}^p \leqslant \|\eta^{(n)}\|_{S_2}^p \leqslant \frac{2^{p-1}\gamma_{2r}^{1-p/2}}{r^{1-p/2}}\|\eta^{(n)}\|_{S_p}^p$$

and also $\|\eta^{(n+1)}\|_{S_2}^p \leqslant \frac{2^{p-1}\gamma_{2r}^{1-p/2}}{r^{1-p/2}}\|\eta^{(n+1)}\|_{S_p}^p$. Inserting that in (66) and dividing by $\|\eta^{(n+1)}\|_{S_p}^p$, we obtain

$$\|\eta^{(n+1)}\|_{S_p}^p \leqslant 2^{4p}(1+\gamma_{2r})^2\gamma_{2r}^{2-p}\Big(\frac{d-r}{r}\Big)^{2-p}\frac{r^{p/2}(s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p\|\eta^{(n)}\|_{S_\infty}^{p(1-p)}\|\eta^{(n)}\|_{S_p}^p.$$

Under the assumption that (45) holds, it follows from this and (65) that

$$\|\eta^{(n+1)}\|_{S_\infty}^p \leqslant \|\eta^{(n)}\|_{S_\infty}^p \text{ and } \|\eta^{(n+1)}\|_{S_p}^p \leqslant \|\eta^{(n)}\|_{S_p}^p \tag{67}$$

for $n=\bar{n}$, which also entails the statement of Theorem 5 for this iteration.

Let now $n' > \bar{n}$ such that (67) is true for all n with $n' > n \geqslant \bar{n}$. If $\sigma_{r+1}(X^{(n')}) \leqslant \epsilon^{(n'-1)}$, then $\epsilon^{(n')} = \sigma_{r+1}(X^{(n')})$ and the arguments from above show (67) also for $n=n'$.

Otherwise $\sigma_{r+1}(X^{(n')}) > \epsilon^{(n'-1)}$ and there exists $n' > n'' \geqslant \bar{n}$ such that $\epsilon^{(n')} = \epsilon^{(n'')} = \sigma_{r+1}(X^{(n'')})$. Then

$$\|\eta^{(n'+1)}\|_{S_\infty}^p \leqslant 2^p\frac{\gamma_{2r}^{2-p}}{r^{2-p}}\Bigg(\sum_{i=r+1}^{d}\Big(\frac{\sigma_i^2(X^{(n')})}{\epsilon^{(n'')2}}+1\Big)^{\frac{p}{2}}\Bigg)^{2-p}\frac{7^p r^{p/2}\max(r,d-r)^{p/2}(s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p\|\eta^{(n')}\|_{S_\infty}^{p(2-p)}$$

and we compute

$$\begin{aligned}&\Bigg(\sum_{i=r+1}^{d}\Big(\frac{\sigma_i^2(X^{(n')})}{\epsilon^{(n'')2}}+1\Big)^{\frac{p}{2}}\Bigg)^{2-p} \leqslant \Bigg(\sum_{i=r+1}^{d}\frac{\sigma_i^p(X^{(n')})}{\epsilon^{(n'')p}}+(d-r)\Bigg)^{2-p} \leqslant \left(\frac{\|\eta^{(n')}\|_{S_p}^p}{\epsilon^{(n'')p}}+(d-r)\right)^{2-p}\\ &\leqslant \left(\frac{\|\eta^{(n'')}\|_{S_p}^p}{\epsilon^{(n'')p}}+(d-r)\right)^{2-p} \leqslant \left(\frac{2(1+\gamma_{2r})\|X_{T_c}^{(n'')}\|_{S_p}^p}{(1-\gamma_{2r})\epsilon^{(n'')p}}+(d-r)\right)^{2-p} \leqslant \left(\frac{3+\gamma_{2r}}{1-\gamma_{2r}}\right)^{2-p}(d-r)^{2-p},\end{aligned}$$

using that X_0 is a matrix of rank at most r in the second inequality, the inductive hypothesis in the third and an analogue of Lemma 9 for a Schatten-p quasinorm on the left hand side (cf. [33, Lemma 3.2] for the corresponding result for $p = 1$). The latter argument uses the assumption on the null space property. This shows that

$$\|\eta^{(n'+1)}\|_{S_\infty}^p \leqslant \mu \|\eta^{(n')}\|_{S_\infty}^{p(2-p)}$$

for

$$\widetilde{\mu} := 2^{4p}\gamma_{2r}^{2-p}\Big(\frac{(3+\gamma_{2r})(d-r)}{(1-\gamma_{2r})r}\Big)^{2-p}\frac{r^{p/2}(s_r^0)^{p(p-1)}}{(1-\rho)^{2p}}\kappa^p \max\Big(2^p(d-r)^{\frac{p}{2}}, (1+\gamma_{2r})^2\Big), \tag{68}$$

and $\|\eta^{(n'+1)}\|_{S_\infty}^p \leqslant \|\eta^{(n')}\|_{S_\infty}^p$ under the assumption (45) of Theorem 5, as $\widetilde{\mu} \leqslant \mu$ with μ as in (46). Indeed $\widetilde{\mu} \leqslant \mu$ since

$$\max\Big(2^p(d-r)^{\frac{p}{2}}, (1+\gamma_{2r})^2\Big)\Big(\frac{d-r}{r}\Big)^{2-p} r^{p/2} \leqslant 2^p(1+\gamma_{2r})^2\Big(\frac{d-r}{r}\Big)^{2-p/2} r^p.$$

The same argument shows that $\|\eta^{(n'+1)}\|_{S_p}^p \leqslant \|\eta^{(n')}\|_{S_p}^p$, which finishes the proof. □

Proof of Theorem 4(iii). The statement follows from Theorem 5, since, for $\widetilde{r} \leqslant r$, an operator Φ fulfilling the Schatten-p NSP of order $2r$ with constant $\gamma_{2r} < 1$ fulfills trivially the Schatten-p NSP of order $2\widetilde{r}$ with constant $\gamma_{2\widetilde{r}} \leqslant \gamma_{2r} < 1$. □

Finally, we show Theorem 2 using Theorem 5.

Proof of Theorem 2 and of Theorem 1(iii). We note that Lemma 8 implies that for a fixed $0 < p \leqslant 1$, there exists a constant $\rho_2 > 1$ and a constant $\tau(p, \rho_2) > 0$ (depending on p and ρ_2) such that if $m \geqslant \rho_2 r(d_1 + d_2)$, the statement of Theorem 2 is fulfilled with the constant μ defined as

$$\mu = 2^{5p}\tau(p,\rho_2)\Big(\frac{d-r}{r}\Big)^{2-\frac{p}{2}} r^p \kappa^p \frac{\sigma_r(X_0)^{p(p-1)}}{(1-\rho)^{2p}} \tag{69}$$

with high probability. This follows since we can define $\tau(p, \rho_2)$ such that

$$\gamma_{2r}^{2-p}\frac{(3+\gamma_{2r})^{2-p}(1+\gamma_{2r})^2}{(1-\gamma_{2r})^{2-p}} \leqslant \tau(p,\rho_2)$$

with high probability. This argument also shows Theorem 1(iii) for $\widetilde{r} = r$. If $\widetilde{r} < r$, define μ as in (69), but replace $\sigma_r(X_0)$ with $\sigma_{\widetilde{r}}(X_0)$ in the formula, and the claim follows analogously. □

Remark 7. *We note that the left- and right-sided weight matrices of previous IRLS approaches [19, 39] at iteration n could be expressed in our notation as*

$$\mathbf{I}_{d_2} \otimes W_L^{(n)} := \mathbf{I}_{d_2} \otimes U^{(n)}(\bar{\Sigma}_{d_1}^{(n)})^{p-2}U^{(n)*}$$

and

$$W_R^{(n)} \otimes \mathbf{I}_{d_1} := V^{(n)}(\bar{\Sigma}_{d_2}^{(n)})^{p-2}V^{(n)*} \otimes \mathbf{I}_{d_1},$$

respectively, cf. section 2.3.

Let

$$T^{(n)} := \operatorname{span}\big\{u_i^{(n)}y^*, xv_i^{(n)*} \mid x \in \mathbb{R}^{d_1}, y \in \mathbb{R}^{d_2}, i \in [r]\big\},$$

where $u_1^{(n)}, \dots, u_r^{(n)}$ resp. $v_1^{(n)}, \dots, v_r^{(n)}$ are the first r left and right singular vectors of $X^{(n)}$. $T^{(n)}$ is a space that can be considered as a generalized support *of the best rank-r approximation of $X^{(n)}$.*

The fact that left- or right-sided weight matrices do not lead to algorithms with superlinear convergence rates for $p < 1$ can be explained by noting that there are always parts of the space $T^{(n)}$ that

are equipped with too large weights if $X^{(n)} = U^{(n)}\Sigma^{(n)}V^{(n)*}$ is already approximately low-rank. In particular, proceeding as in (58)*, we obtain for* $\mathbf{I}_{d_2} \otimes W_L^{(n)}$

$$\begin{aligned}\left\| \left[\left(\mathbf{I}_{d_2} \otimes W_L^{(n)} \right) (X_0)_{\mathrm{vec}} \right]_{\mathrm{mat}} \right\|_{S_2}^2 &= \left\| \left(\bar{\Sigma}_T^{(n)} \right)^{p-2} U_T^{(n)*} X_0 V_T^{(n)} \right\|_{S_2}^2 + \left\| \left(\bar{\Sigma}_T^{(n)} \right)^{p-2} U_T^{(n)*} X_0 V_{T_c}^{(n)} \right\|_{S_2}^2 \\ &\quad + \left\| \left(\bar{\Sigma}_{T_c}^{(n)} \right)^{p-2} U_{T_c}^{(n)*} X_0 V_T^{(n)} \right\|_{S_2}^2 + \left\| \left(\bar{\Sigma}_{T_c}^{(n)} \right)^{p-2} U_{T_c}^{(n)*} X_0 V_{T_c}^{(n)} \right\|_{S_2}^2\end{aligned}$$

if $\bar{\Sigma}_T^{(n)}$ denotes the diagonal matrix with the first r non-zero entries of $\bar{\Sigma}_{d_1}^{(n)}$ and $\bar{\Sigma}_{T_c}^{(n)}$ the one of the remaining entries.

Here, the third of the four summands would become too large for $p < 1$ to allow for a superlinear convergence when the last $d - r$ singular values of $X^{(n)}$ approach zero. An analogous argument can be used for the right-sided weight matrix $W_R^{(n)} \otimes \mathbf{I}_{d_1}$ and, notably, also for arithmetic mean weight matrices $W_{(arith)}^{(n)} = \mathbf{I}_{d_2} \otimes W_L^{(n)} + W_R^{(n)} \otimes \mathbf{I}_{d_1}$, cf. section 2.4.

7 Acknowledgements

The two authors acknowledge the support and hospitality of the Hausdorff Research Institute for Mathematics (HIM) during the early stage of this work within the HIM Trimester Program "Mathematics of Signal Processing". C.K. acknowledges the support by the German Research Foundation (DFG) in the context of the Emmy Noether Junior Research Group "Randomized Sensing and Quantization of Signals and Images" (KR 4512/1-1). J.S. acknowledges the support by the DFG through the D-A-CH project no. I1669-N26 and through the international research training group IGDK 1754 "Optimization and Numerical Analysis for Partial Differential Equations with Nonsmooth Structures". The authors thank Ke Wei for providing code of his implementations. They also thank Massimo Fornasier for helpful discussions.

References

[1] A. Ahmed and J. Romberg. Compressive multiplexing of correlated signals. *IEEE Transactions on Inf. Theory*, 61(1):479–498, 2015.

[2] K. M. R. Audenaert. A generalisation of Mirsky's singular value inequalities. ArXiv preprint:1410.4941, 2014.

[3] D. S. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas (Second Edition)*. Princeton University Press, 2009.

[4] J. D. Blanchard, J. Tanner, and K. Wei. CGIHT: conjugate gradient iterative hard thresholding for compressed sensing and matrix completion. *Inf. Inference*, 4(4):289–327, 2015.

[5] E. J. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM J. on Imaging Sci.*, 6(1):199–225, 2013.

[6] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Inf. Theory*, 61(4):1985–2007, 2015.

[7] E. J. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Inf. Theory*, 57(4):2342–2359, April 2011.

[8] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations Comput. Math.*, 9(6):717–772, 2009.

[9] E. J. Candès, T. Strohmer, and V. Voroninski. PhaseLift: Exact and Stable Signal Recovery from Magnitude Measurements via Convex Programming. *Commun. on Pure Appl. Math.*, 66(8):1241–1274, 2013.

[10] R. Chartrand. Exact reconstructions of sparse signals via nonconvex minimization. *IEEE Signal Process. Lett.*, 14:707–710, 2007.

[11] J. A. Chavez-Dominguez and D. Kutzarova. Stability of low-rank matrix recovery and its connections to Banach space geometry. *J. Math. Analysis Appl.*, 427(1):320 – 335, 2015.

[12] B. Dacorogna. *Direct Methods in the Calculus of Variations*. Springer-Verlag New York, Inc., 1989.

[13] I. Daubechies, R. DeVore, M. Fornasier, and C. Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Commun. on Pure Appl. Math.*, 63:1–38, 2010.

[14] J. Davenport, Mark A.; Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE J. Sel. Top. Signal Process.*, 10:608–622, 06 2016.

[15] Y. Eldar, D. Needell, and Y. Plan. Uniqueness conditions for low-rank matrix recovery. *Appl. Comput. Harmon. Anal.*, 33(2):309–314, 2012.

[16] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Electrical Engineering Department Stanford University, 2002.

[17] M. Fornasier. Numerical methods for sparse recovery. In M. Fornasier, editor, *Theoretical Foundations and Numerical Methods for Sparse Recovery*, volume 9 of *Radon Series on Computational and Applied Mathematics*, pages 93–200. De Gruyter, Berlin, 2010.

[18] M. Fornasier, S. Peter, H. Rauhut, and S. Worm. Conjugate gradient acceleration of iteratively re-weighted least squares methods. *Comput. Optim. Appl.*, pages 1–55, 2016.

[19] M. Fornasier, H. Rauhut, and R. Ward. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM J. on Optim.*, 21(4):1614–1640, 2011. [using `IRLS-FRW` ('IRLS-M') algorithm, code from `https://github.com/rward314/IRLSM`].

[20] S. Foucart. Concave Mirsky inequality and low-rank recovery. `http://www.math.tamu.edu/~foucart/publi/Mirsky.pdf`, preprint, 2016.

[21] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Springer New York, 2013.

[22] Y. Gao, J. Peng, S. Yue, and Y. Zhao. On the null space property of ℓ_q-minimization for $0 < q \leqslant 1$ in compressed sensing. *J. Funct. Spaces*, 2015:4203–4215, 2015.

[23] I. Gohberg, S. Goldberg, and N. Krupnik. *Traces and determinants of linear operators*, volume 116. Springer, 2012.

[24] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35(12):61–70, 1992.

[25] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Inf. Theory*, 57(3):1548–1566, 2011.

[26] D. Gross, F. Krahmer, and R. Kueng. A partial derandomization of phaselift using spherical designs. *J. Fourier Analysis Appl.*, 21(2):229–266, 2015.

[27] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105:150401, 2010.

[28] J. P. Haldar and D. Hernando. Rank-constrained solutions to linear matrix equations using powerfactorization. *IEEE Signal Process. Lett.*, 16(7):584–587, July 2009. [using `p_MC_AltMin` (Alternating Minimization) algorithm, partial code available at `http://mr.usc.edu/download/irpf/`].

[29] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

[30] P. Jain, R. Meka, and I. S. Dhillon. Guaranteed rank minimization via singular value projection. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 937–945. Curran Associates, Inc., 2010.

[31] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 665–674, New York, NY, USA, 2013. ACM.

[32] A. Jameson. Solution of the equation ax + xb = c by inversion of an m x m or n x n matrix. *SIAM J. on Appl. Math.*, 16(5):1020–1023, 1968.

[33] M. Kabanava, R. Kueng, H. Rauhut, and U. Terstiege. Stable low-rank matrix recovery via null space properties. *Inf. Inference: A J. IMA*, 5(4):405, 2016.

[34] A. Kyrillidis and V. Cevher. Matrix recipes for hard thresholding methods. *J. Math. Imaging Vis.*, 48(2):235–265, 2014. [using `Matrix ALPS II` ('Matrix ALgrebraic PursuitS II') algorithm, code from `http://akyrillidis.github.io/projects/`].

[35] Z. Liu, A. Hansson, and L. Vandenberghe. Nuclear norm system identification with missing inputs and outputs. *Syst. & Control. Lett.*, 62(8):605 – 612, 2013.

[36] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM J. on Matrix Analysis Appl.*, 31(3):1235–1256, 2010.

[37] J. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley Series in Probability and Statistics: Texts and References Section. Wiley, 1999.

[38] B. Mishra, G. Meyer, F. Bach, and R. Sepulchre. Low-rank optimization with trace norm penalty. *SIAM J. on Optim.*, 23(4):2124–2149, 2013.

[39] K. Mohan and M. Fazel. Iterative reweighted algorithms for matrix rank minimization. *The J. Mach. Learn. Res.*, 13(1):3441–3473, 2012. [using `IRLS-MF` ('IRLS-p') algorithm, code from `https://faculty.washington.edu/mfazel/`].

[40] S. Oymak, K. Mohan, M. Fazel, and B. Hassibi. A simplified approach to recovery conditions for low rank matrices. In *2011 IEEE International Symposium on Information Theory Proceedings, ISIT 2011, St. Petersburg, Russia*, pages 2318–2322, 2011.

[41] D. Park, A. Kyrillidis, C. Caramanis, and S. Sanghavi. Finding Low-rank Solutions to Matrix Problems, Efficiently and Provably. *ArXiv preprints:1606.03168*, June 2016. [using `BFGD` ('Bi-Factored Gradient Descent') algorithm, code from `http://akyrillidis.github.io/projects/`].

[42] D. Pimentel-Alarcón, N. Boston, and R. D. Nowak. A Characterization of Deterministic Sampling Patterns for Low-Rank Matrix Completion. *ArXiv preprints:1503.02596*, 2015.

[43] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, 2010.

[44] B. Recht, W. Xu, and B. Hassibi. Null space conditions and thresholds for rank minimization. *Math. Program.*, 127(1):175–202, 2011.

[45] E. Schost and P.-J. Spaenlehauer. A quadratically convergent algorithm for structured low-rank approximation. *Foundations Comput. Math.*, 16(2):457–492, 2016.

[46] N. Srebro, J. Rennie, and T. S. Jaakkola. Maximum-margin matrix factorization. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1329–1336. MIT Press, 2005.

[47] M. Stewart. Perturbation of the SVD in the presence of small singular values. *Linear Algebr. its Appl.*, 419(1):53 – 77, 2006.

[48] J. Sun, Q. Qu, and J. Wright. When are nonconvex problems not scary? ArXiv preprint:1510.06096, 2015.

[49] R. Sun and Z. Q. Luo. Guaranteed matrix completion via nonconvex factorization. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 270–289, 2015.

[50] J. Tanner and K. Wei. Normalized Iterative Hard Thresholding for Matrix Completion. *SIAM J. on Sci. Comput.*, 35(5):S104–S125, 2013.

[51] J. Tanner and K. Wei. Low rank matrix completion by alternating steepest descent methods. *Appl. Comput. Harmon. Analysis*, 40(2):417 – 429, 2016. [using `ASD` ('Alternating Steepest Descent') algorithm, code from `https://www.math.ucdavis.edu/~kewei/publications.html` resp. `https://www.math.ucdavis.edu/~kewei/code/mc20140528.tar`].

[52] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht. Low-rank Solutions of Linear Matrix Equations via Procrustes Flow. *ArXiv preprint: 1507.03566*, 2015.

[53] C. F. Van Loan. The ubiquitous Kronecker product. *J. Comput. Appl. Math.*, 123(12):85 – 100, 2000.

[54] B. Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM J. on Optim.*, 23(2):1214–1236, 2013. [using `Riemann_Opt` ('Riemannian Optimization') algorithm, code from `http://www.unige.ch/math/vandereycken/matrix_completion.html`].

[55] P.-Å. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numer. Math.*, 12(1):99–111, 1972.

[56] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung. Guarantees of Riemannian Optimization for Low Rank Matrix Completion. *ArXiv preprint:1603.06610*, 2016.

[57] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung. Guarantees of Riemannian Optimization for Low Rank Matrix Recovery. *SIAM J. Matrix Anal. Appl.*, 37(3):1198–1222, 2016.

[58] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Math. Program. Comput.*, 4(4):333–361, 2012.

[59] Q. Zheng and J. Lafferty. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 109–117. Curran Associates, Inc., 2015.

[60] Q. Zheng and J. Lafferty. Convergence Analysis for Rectangular Matrix Completion Using Burer-Monteiro Factorization and Gradient Descent. *ArXiv preprint:1605.07051*, 2016.

A Appendix: Proof of Lemma 4: Critical points of $f^p_{X,\epsilon}$

Let us without loss of generality consider the case $d = d_1 = d_2$ and define the set

$$\Omega = \{W \in M_{d\times d} \text{ s.t. } \operatorname{rank}(W) = d\}\,.$$

As already mentioned in (26), the harmonic mean matrix $\widetilde{W}(W)$ can be then rewritten as

$$\widetilde{W}(W) = 2\widetilde{W}_1\big(\widetilde{W}_1 + \widetilde{W}_2\big)^{-1}\widetilde{W}_2 = 2(\widetilde{W}_1^{-1} + \widetilde{W}_2^{-1})^{-1}$$

for $W \in \Omega$ with the definitions $\widetilde{W}_1 := \mathbf{I}_d \otimes (WW^*)^{\frac{1}{2}}$ and $\widetilde{W}_2 = (W^*W)^{\frac{1}{2}} \otimes \mathbf{I}_d$. For $W \in \Omega$, we reformulate the auxiliary functional such that

$$\begin{aligned} f^p_{X,\epsilon}(W) = \mathcal{J}^p(X,\epsilon,W) &= \frac{p}{2}\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} + \frac{\epsilon^2 p}{2}\sum_{i=1}^d \sigma_i(W) + \frac{2-p}{2}\sum_{i=1}^d \sigma_i(W)^{\frac{p}{(p-2)}} \\ &= \frac{p}{2}\|X_{\mathrm{vec}}\|^2_{\ell_2(\widetilde{W}(W))} + \frac{\epsilon^2 p}{2}\|(W^*W)^{1/2}\|_F^2 + \frac{2-p}{2}\|(W^*W)^{\frac{p}{2(p-2)}}\|_F^2. \end{aligned}$$

To identify the set of critical points of $f^p_{X,\epsilon}(W)$ located in Ω, we compute its derivative with respect to W using the derivative rules (7), (12), (13), (15), (16), (18), (20) in Chapter 8.2 and Theorem 3 in Chapter 8.4 of [37] in the following. Using the notation of [37], we calculate

$$\begin{aligned} \frac{\partial f^p_{X,\epsilon}(W)}{\partial W} = -\frac{p}{2}\operatorname{tr}\left(X^*_{\mathrm{vec}}\widetilde{W}\frac{\partial\widetilde{W}^{-1}}{\partial W}\widetilde{W}X_{\mathrm{vec}}\right) + \frac{p\epsilon^2}{4}\left(\operatorname{tr}\left(W(W^*W)^{-\frac{1}{2}}\partial W^*\right) + \operatorname{tr}((W^*W)^{-\frac{1}{2}}W^*\partial W)\right) \\ -\frac{p}{4}\left(\operatorname{tr}\left(W(W^*W)^{\frac{4-p}{2(p-2)}}\partial W^*\right) + \operatorname{tr}((W^*W)^{\frac{4-p}{2(p-2)}}W^*\partial W)\right) \end{aligned}$$

where

$$\begin{aligned} \frac{\partial\widetilde{W}^{-1}}{\partial W} = \frac{1}{2}\frac{\partial\left[(WW^*)^{-\frac{1}{2}}\oplus(W^*W)^{-\frac{1}{2}}\right]}{\partial W} = -\frac{1}{4}\left[\left((W^*W)^{-\frac{3}{2}}W^*\partial W + \partial W^*W(W^*W)^{-\frac{3}{2}}\right)\otimes\mathbf{I}_{d_1}\right] \\ -\frac{1}{4}\left[\mathbf{I}_{d_2}\otimes\left(\partial W(WW^*)^{-\frac{3}{2}}W^* + (WW^*)^{-\frac{3}{2}}W\partial W^*\right)\right]. \end{aligned} \tag{70}$$

We can reformulate the first term as follows using the cyclicity of the trace,

$$\begin{aligned} -\frac{p}{2}\operatorname{tr}\left(X^*_{\mathrm{vec}}\widetilde{W}\frac{\partial\widetilde{W}^{-1}}{\partial W}\widetilde{W}X_{\mathrm{vec}}\right) = \frac{p}{8}\Big[&\operatorname{tr}\left((\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(W^*W)^{-\frac{3}{2}}W^*\partial W\right) \\ &+ \operatorname{tr}\left(W(W^*W)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}\partial W^*\right) \\ &+ \operatorname{tr}\left(W^*(WW^*)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}\partial W\right) \\ &+ \operatorname{tr}\left((\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(WW^*)^{-\frac{3}{2}}W\partial W^*\right)\Big]. \end{aligned}$$

Summarizing the calculations above, we have

$$\begin{aligned}\frac{\partial f^p_{X,\epsilon}(W)}{\partial W} &= \frac{p}{8}\Big[\operatorname{tr}\Big((\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(W^*W)^{-\frac{3}{2}}W^*\partial W\Big)\\ &+ \operatorname{tr}\Big(W(W^*W)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}\partial W^*\Big)\\ &+ \operatorname{tr}\Big(W^*(WW^*)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}\partial W\Big)\\ &+ \operatorname{tr}\Big((\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(WW^*)^{-\frac{3}{2}}W\partial W^*\Big)\Big]\\ &+ \frac{p\epsilon^2}{4}\Big(\operatorname{tr}\Big(W(W^*W)^{-\frac{1}{2}}\partial W^*\Big) + \operatorname{tr}((W^*W)^{-\frac{1}{2}}W^*\partial W)\Big)\\ &- \frac{p}{4}\Big(\operatorname{tr}\Big(W(W^*W)^{\frac{4-p}{2(p-2)}}\partial W^*\Big) + \operatorname{tr}((W^*W)^{\frac{4-p}{2(p-2)}}W^*\partial W)\Big).\end{aligned}$$

To determine the critical points of $f^p_{X,\epsilon}(W)$, we rearrange the terms and equate the derivative with zero, such that

$$\begin{aligned}\frac{\partial f^p_{X,\epsilon}(W)}{\partial W} &= \frac{p}{8}\operatorname{tr}\Big(\Big[(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(W^*W)^{-\frac{3}{2}}W^* + W^*(WW^*)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}\\ &\qquad +2\epsilon^2(W^*W)^{-\frac{1}{2}}W^* - 2(W^*W)^{\frac{4-p}{2(p-2)}}W^*\Big]\partial W\Big)\\ &\quad \frac{p}{8}\operatorname{tr}\Big(\Big[W(W^*W)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}} + (\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(WW^*)^{-\frac{3}{2}}W\\ &\qquad +2\epsilon^2W(W^*W)^{-\frac{1}{2}} - 2W(W^*W)^{\frac{4-p}{2(p-2)}}\Big]\partial W^*\Big)\\ &:= \frac{p}{8}\operatorname{tr}(A\partial W) + \frac{p}{8}\operatorname{tr}(A^*\partial W^*)\\ &= \frac{p}{8}\operatorname{tr}((A\oplus A)\partial W)\\ &= 0,\end{aligned}$$

where

$$\begin{aligned}A = \Big[&(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(W^*W)^{-\frac{3}{2}}W^* + W^*(WW^*)^{-\frac{3}{2}}(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}}(\widetilde{W}X_{\mathrm{vec}})^*_{\mathrm{mat}}\\ &+2\epsilon^2(W^*W)^{-\frac{1}{2}}W^* - 2(W^*W)^{\frac{4-p}{2(p-2)}}W^*\Big].\end{aligned} \tag{71}$$

Now we have to find W such that $A\oplus A = 0$. This implies that all eigenvalues of $A\oplus A = A\otimes\mathbf{I}_d + \mathbf{I}_d\otimes A$ have to be zero. The eigenvalues of the Kronecker sum of two matrices A_1 and A_2 with eigenvalues λ_s and μ_t with $s,t\in[d]$ are the sum of the eigenvalues $\lambda_s+\mu_t$. As in our case $A = A_1 = A_2$ this means that all eigenvalues of A itself have to be zero. This is only possible if A is the zero matrix.

Let $W = U\Sigma V^*\in M_{d\times d}$ with $U\in U_d, V\in U_d, \Sigma\in M_{d\times d}$, where $\Sigma = \operatorname{diag}(\sigma)$ is a diagonal matrix with *ascending* entries. We define the matrix $H = H_{i,j} = \frac{2}{\sigma_i^{-1}+\sigma_j^{-1}}$ for $i = 1,\dots,d, j = 1,\dots,d$ corresponding to the result of reshaping the diagonal of $2(\Sigma\oplus\Sigma)$ into a $d\times d$-matrix. Using (4), we can express $(\widetilde{W}X_{\mathrm{vec}})_{\mathrm{mat}} = U\big(H\circ(U^*XV)\big)V^*$ and denote $B := H\circ(U^*XV)$.

Plugging the decomposition $W = U\Sigma V^*$ into (71), we can therefore calculate that

$$\begin{aligned}\mathrm{A} = 0 &\Leftrightarrow (UBV^*)^*(UBV^*)(V\Sigma^2V^*)^{-3/2}(U\Sigma V^*)^* + (U\Sigma V^*)^*(U\Sigma^2U^*)^*)^{-3/2}(UBV^*)(UBV^*)^*\\ &\quad + 2\epsilon^2(V\Sigma^2V^*)^{-1/2}(U\Sigma V^*)^* - 2(V\Sigma^2V^*)^{\frac{4-p}{2(p-2)}}(U\Sigma V^*)^* = 0\\ &\Leftrightarrow VB^*B\Sigma^{-2}U^* + V\Sigma^{-2}BB^*U^* + 2\epsilon^2V\mathbf{I}_dU^* - 2V\Sigma^{\frac{2}{p-2}}U^* = 0\\ &\Leftrightarrow B^*B\Sigma^{-2} + \Sigma^{-2}BB^* + 2\epsilon^2\mathbf{I}_d - 2\Sigma^{\frac{2}{p-2}} = 0.\end{aligned} \tag{72}$$

We now note that $2\epsilon^2\mathbf{I}_d - 2\Sigma^{\frac{2}{p-2}}$ is diagonal and therefore, $B^*B\Sigma^{-2} + \Sigma^{-2}BB^*$ is diagonal as well. Moreover, observe that $B^*B + \Sigma^{-2}BB^*\Sigma^2$ is again a diagonal matrix and has a symmetric first summand B^*B. As the sum or difference of symmetric matrices is again symmetric also the second summand $\Sigma^{-2}BB^*\Sigma^2$ has to be symmetric, i.e., $\Sigma^{-2}BB^*\Sigma^2 = (\Sigma^{-2}BB^*\Sigma^2)^* = \Sigma^2BB^*\Sigma^{-2}$. We conclude that it has to hold that $BB^*\Sigma^4 = \Sigma^4BB^*$ and hence Σ^4 and BB^* commute.

This is only possible if either Σ is a multiple of the identity or if BB^* is diagonal. Assuming the first case, (72) would imply that also BB^* and B^*B have to be a multiple of the identity. Therefore, this first case, where Σ is a multiple of the identity is a special case of the second possible scenario, where BB^* is diagonal. Hence, it suffices to further consider the more general second case. (Considerations for B^*B can be carried out analogously.)

Diagonality of BB^* only occurs if B is either orthonormal or diagonal. Assuming orthonormality would lead to contradictions with the equations in (72). Hence $B = H \circ (U^*XV)$ can only be diagonal.

Let now be $X = \bar{U}\bar{S}\bar{V}^*$ the singular value decomposition of X. As H has no zero entries due to the full rank of W, this implies the diagonality of $U^*\bar{U}\bar{S}\bar{V}^*V$. Consequently, U and V can only be chosen such that $P = [U^*\bar{U}]_{d\times d}$ and $P^* = [\bar{V}^*V]_{d\times d}$ for a permutation matrix $P \in U_d$. The reshuffled indexing corresponding to P is denoted by $p(i) \in [d]$ for $i \in [d]$. Having in mind that $H_{ii} = \sigma_i$ for $i \in [d]$, we obtain

$$
\begin{aligned}
&(H \circ (P\bar{S}P^*))^*(H \circ (P\bar{S}P^*))\Sigma^{-2} + \Sigma^{-2}(H \circ (P\bar{S}P^*))(H \circ (P\bar{S}P^*))^* + 2\epsilon^2\mathbf{I}_d - 2\Sigma^{\frac{2}{p-2}} = 0 \\
\Leftrightarrow\quad & 2\bar{s}_{p(i)}^2 + 2\epsilon^2 = 2\sigma_i^{\frac{2}{p-2}} \text{ for all } i \in [d] \\
\Leftrightarrow\quad & \sigma_i = (\bar{s}_{p(i)}^2 + \epsilon^2)^{\frac{p-2}{2}} \text{ for all } i \in [d].
\end{aligned}
$$

As the diagonal of Σ as assumed to have ascending entries and the diagonal of $\bar{S}$ has descending entries, the permutation matrix P has to be equal to the identity matrix. From $P = \mathbf{I}_d$, it follows that $U = \bar{U}$ and $V = \bar{V}$ and hence $\Sigma = (\bar{S}^2 + \epsilon^2\mathbf{I}_d)^{\frac{p-2}{2}}$.

We summarize our calculations by stating that

$$
W_{\text{opt}} = \bar{U}\Sigma\bar{V}^* = \bar{U}(\bar{S}^2 + \epsilon^2\mathbf{I}_d)^{\frac{p-2}{2}}\bar{V}^*
$$

is the only critical point of $f_{X,\epsilon}^p$ on the domain Ω.

The results extend for the case $d_1 \neq d_2$, where the definition of $\widetilde{W}(W)$ is adapted by introducing the Moore-Penrose pseudo inverse of $(WW^*)^{1/2}$

$$
\widetilde{W}(W) = 2\widetilde{W}_1\big(\widetilde{W}_1 + \widetilde{W}_2\big)^{-1}\widetilde{W}_2 = 2(\widetilde{W}_1^{+} + \widetilde{W}_2^{-1})^{-1}.
$$

The corresponding derivative rule as pointed out in Theorem 5 in Chapter 8.4 of [37] can be used for the calculation in (70).